

Security and Cryptography 2017

Mirosław Kutylowski

grading criteria: 50% exam, 50% assignments

skills to be learned: developing end-to-end security systems, they must be flawless!

rules: do not memorize the standards, they change. Only the skills are important

presence: obligatory during the lectures

exam date: TBA (early to enable internships in February), optionally: midterm exam(s)

I. EXAMPLE TO LEARN FROM: PKI FAILURE

PKI - Public Key Infrastructure

- strong authentication of digital documents with digital signatures seems to be possible
- in fact we get an evidence that the holder of a private key has created a signature
- who holds the key? PKI has to provide a certified answer to this question
- PKI is not a cryptographic solution - it is an organizational framework (using some crypto tools)

PKI, X.509 standard

- a certificate binds a public key with an ID of its alleged owner,
- a couple of other fields, like validity date, key usage, certification policy, ...
- certificate signed by CA (Certification Authority)
- tree of CA's (or directed acyclic graph), with roots as "root of trust"
- status of a certificate may change - revocation
- checking status methods: CRL, OCSP

reasons for PKI failure:

a nice concept of digital signatures but

1. big infrastructure required:
 - big effort and cost

- long time planning needed (so possible in China, but not in Europe)
 - unclear financial return
2. scope of necessary coordination,
 - in order to work must be designed at least for the Common Market
 - example of killing the concept: link to certification policy in Polish
 3. lack of interoperability (sometimes as business goal)
 - companies make efforts to eliminate competition
 - standardization may be focused on having market shares
 4. necessary trust in roots
 - how do you know that the root is honest?
 5. registration: single point of fraud, (e.g. with fake breeding documents)
 - once you get a certificate you may forge signatures
 6. responsibility of CA
 - financial risk – based on risk or responsibility
 7. cost - who will pay? For the end user the initial cost is too high.
 - certificates are too expensive for just a few signatures (at least initially)
 8. legal strength of signatures
 - if scheme broken or signing devices turn out to be insecure you are anyway responsible for the signatures
 9. unsolved problem of revocation: possible to check the status in the past but not now

reason: mismatch of requirements and interests with the designed solution

MAJOR PROBLEM: how to design/buy good systems?

II. COMMON CRITERIA FRAMEWORK

<http://www.commoncriteriaportal.org/>

Problem: somebody has to deploy a secure IT system, how to purchase it?

- problematic requirements according to BSI guide:
 - i. **incomplete** – forgetting some threats is common
 - ii. **not embedded:** not corresponding really to the environment where the product has to be deployed
 - iii. **implicit:** customer has in mind but the developer might be unaware of them

- iv. **not testable:** ambiguous, source of legal disputes, ...
 - v. **too detailed:** unnecessary details make it harder to adjust the design
 - vi. **unspecified meaning:** e.g. “*protect privacy*”
 - vii. **inconsistent:** e.g. ignoring trade-offs
- *specification-based purchasing process* versus *selection-based purchasing process*
 - the user is not capable of determining the properties of the product himself: too complicated, too specialized knowledge required, a single error makes the product useless
 - specifications of concrete products might be useless for the customers – hard to understand and compare the products
 - informal specifications and descriptions, no crucial data

Idea of Common Criteria Framework:

- standardize the process of
 - designing requirements (Protection Profile, PP) (customer)
 - designing products (Security Target ST), (developer)
 - evaluation of products (licensed labs checking conformance of implementation with the documentation) (certification body)
- international agreement of bodies from some countries (USA, France, UK, Germany, India, Turkey, Sweden, Spain, Australia, Canada, Malaysia, Netherlands, Korea, New Zeland, Italy, Turkey) but Israel only “consuming”, no Poland, China, Singapore,
- idea: ease the process, reuse work, build up from standard components
- typically ST as a response for PP:
 - more detailed
 - maybe chooses some concrete options
 - maybe fulfills more requirements (more PP)
 - relation with PP should be testable

Value:

- CC certification does not mean a product is secure
- it only says that it has been developed according to PP
- assurance level concerns only the stated requirements , e.g. trivial requirements ⇒ high EAL level (common mistake in public procurement: EAL level ... without specifying PP)

- but it is cleaning up the zoo of different assumptions, descriptions, ...

Example for PP: BAC (Basic Access Control)

- used to secure wireless communication between a reader and an e-Passport (of an old generation)
- encryption primitive

$$EM(K, S) = \text{Enc}(KB_{\text{Enc}}, S) \parallel \text{MAC}(KB_{\text{Mac}}, \text{Enc}(KB_{\text{Enc}}, S), S)$$

where the key K is $(KB_{\text{Enc}}, KB_{\text{Mac}})$

- steps:
 1. The MRTD chip sends a nonce r_{PICC} to the terminal
 2. The terminal sends the encrypted challenge

$$e_{\text{PCD}} = EM(K, r_{\text{PCD}}, r_{\text{PICC}}, K_{\text{PCD}})$$

to the MRTD chip, where r_{PICC} is the MRTD chip's nonce, r_{PCD} is the terminal's randomly chosen nonce, and K_{PCD} is keying material for the generation of the session keys.

3. The MRTD chip decrypts and verifies r_{PICC} , responds with

$$e_{\text{PICC}} = EM(K, r_{\text{PICC}}, r_{\text{PCD}}, K_{\text{PICC}})$$

4. The terminal decrypts and verifies r_{PCD}
5. both sides derive $K_{\text{Enc}}, K_{\text{Mac}}$ from the master key

$$K_{\text{PICC}} \text{ XOR } K_{\text{PCD}}$$

and a sequence number derived from the random nonces (key derivation function)

- **K derived from information available on the machine readable zone (optical reader applied, not available via wireless connection)**
- implementation: biometric passports.
- a simple system. Really?

Common Criteria Protection Profile Machine Readable Travel Document with ICAO Application, Basic Access Control BSI-CC-PP-0055

1. Introduction

aimed for customers looking for proper products, overview

1.1 PP reference

basic data, registration data

Title: Protection Profile - Machine Readable Travel Document with ICAO Application and Basic Access Control (MRTD-PP)

Sponsor: Bundesamt für Sicherheit in der Informationstechnik CC Version: 3.1 (Revision 2)

Assurance Level: The minimum assurance level for this PP is EAL₄ augmented.

General Status: Final

Version Number: 1.10

Registration: BSI-CC-PP-0055

Keywords: ICAO, machine readable travel document, basic access control

1.2 TOE Overview

- Target of Evaluation
- "is aimed at potential consumers who are looking through lists of evaluated TOEs/Products to find TOEs that may meet their security needs, and are supported by their hardware, software and firmware"
- important sections:
 - Usage and major security features of the TOE
 - TOE type
 - Required non-TOE hardware/software/firmware
- Definition, Type
which parts, which general purpose, which functionalities are present and which are missing, e.g. ATM card with no contactless payments
- Usage and security features
crucial properties of the system (high level) and security features from the point of view of the security effect and not how it is achieved
- life cycle
the product in the whole life cycle including manufacturing, delivery and destroying
- Required non-TOE hardware/software/firmware: other components that can be crucial for evaluation

2. Conformance Claim

- CC Conformance Claim: version of CC
- PP claim: other PP taken into account in a plug-and-play way
- Package claim: which EAL package level

EAL packages:

- The CC formalizes assurance into 6 categories (the so-called "assurance classes" which are further subdivided into 27 sub-categories (the so-called "assurance families"). In each assurance family, the CC allows grading of an evaluation with respect to that assurance family.
- 7 predefined ratings, called evaluation assurance levels or EALs. called EAL1 to EAL7, with EAL1 the lowest and EAL7 the highest
- Each EAL can be seen as a set of 27 numbers, one for each assurance family. EAL1 assigns a rating of 1 to 13 of the assurance families, and 0 to the other 14 assurance families, while EAL2 assigns the rating 2 to 7 assurance families, the rating 1 to 11 assurance families, and 0 to the other 9 assurance families
- monotonic: EAL_{n+1} gives at least the same assurance level as EAL_n in each assurance families
- levels:
 - EAL1: Functionally Tested:
 - correct operation, no serious threats
 - minimal effort from the manufacturer
 - EAL2: Structurally Tested
 - delivery of design information and test results,
 - effort on the part of the developer than is consistent with good commercial practice.
 - EAL3: Methodically Tested and Checked
 - maximum assurance from positive security engineering at the design stage without substantial alteration of existing sound development practices.
 - developers or users require a moderate level of independently assured security, and require a thorough investigation of the TOE and its development without substantial re-engineering.
 - EAL4: Methodically Designed, Tested and Reviewed
 - maximum assurance from positive security engineering based on good commercial development practices which, though rigorous, do not require substantial specialist knowledge, skills, and other resources.
 - the highest level at which it is likely to be economically feasible to retrofit to an existing product line.
 - EAL5: Semiformally Designed and Tested
 - EAL6: Semiformally Verified Design and Tested
 - EAL7: Formally Verified Design and Tested

- assurance classes:
 - development:
 - ADV_ARC - 1 1 1 1 1 1 architecture requirements
 - ADV_FSP 1 2 3 4 5 5 6 functional specifications
 - ADV_IMP - - - 1 1 2 2 implementation representation
 - ADV_INT - - - - 2 3 3 “is designed and structured such that the likelihood of flaws is reduced and that maintenance can be more readily performed without the introduction of flaws”?
 - ADV_SPM - - - - - 1 1 security policy modeling
 - ADV_TDS - 1 2 3 4 5 6 TOE design
 - guidance documents
 - AGD_OPE 1 1 1 1 1 1 1 Operational user guidance
 - AGD_PRE 1 1 1 1 1 1 1 Preparative procedures
 - life-cycle support
 - ALC_CMC 1 2 3 4 4 5 5 CM capabilities
 - ALC_CMS 1 2 3 4 5 5 5 CM scope
 - ALC_DEL - 1 1 1 1 1 1 1 Delivery
 - ALC_DVS - - 1 1 1 2 2 Development security
 - ALC_FLR - - - - - - - Flaw remediation
 - ALC_LCD - - 1 1 1 1 2 Life-cycle definition
 - ALC_TAT - - - 1 2 3 3 Tools and techniques
 - security target evaluation
 - ASE_CCL 1 1 1 1 1 1 1 1 Conformance claims
 - ASE_ECD 1 1 1 1 1 1 1 1 Extended components definition
 - ASE_INT 1 1 1 1 1 1 1 1 ST introduction
 - ASE_OBJ 1 2 2 2 2 2 2 2 Security objectives
 - ASE_REQ 1 2 2 2 2 2 2 2 Security requirements
 - ASE_SPD - 1 1 1 1 1 1 1 Security problem definition
 - ASE_TSS - 1 1 1 1 1 1 1 TOE summary specification
 - tests
 - ATE_COV 1 2 2 2 3 3 Coverage

- ATE_DPT 1 1 3 3 4 Depth
 - ATE_FUN 1 1 1 1 2 2 Functional tests
 - ATE_IND 1 2 2 2 2 2 3 Independent testing
- vulnerability assesment
- AVA_VAN 1 2 2 3 4 5 5 Vulnerability analysis
- for example, a product could score in the assurance family developer test coverage (ATE_COV):
 - 0: It is not known whether the developer has performed tests on the product;
 - 1: The developer has performed some tests on some interfaces of the product;
 - 2: The developer has performed some tests on all interfaces of the product;
 - 3: The developer has performed a very large amount of tests on all interfaces of the product
 - example more formal: ALC_FLR
 - ALC_FLR.1:
 - *The flaw remediation procedures documentation shall describe the procedures used to track all reported security flaws in each release of the TOE.*
 - *The flaw remediation procedures shall require that a description of the nature and effect of each security flaw be provided, as well as the status of finding a correction to that flaw.*
 - *The flaw remediation procedures shall require that corrective actions be identified for each of the security flaws.*
 - *The flaw remediation procedures documentation shall describe the methods used to provide flaw information, corrections and guidance on corrective actions to TOE users.*
 - ALC_FLR.2:
 - ALC_FLR.1 as before
 - *The flaw remediation procedures shall describe a means by which the developer receives from TOE users reports and enquiries of suspected security flaws in the TOE.*
 - *The procedures for processing reported security flaws shall ensure that any reported flaws are remediated and the remediation procedures issued to TOE users.*
 - *The procedures for processing reported security flaws shall provide safeguards that any corrections to these security flaws do not introduce any new flaws.*
 - *The flaw remediation guidance shall describe a means by which TOE users report to the developer any suspected security flaws in the TOE.*

- ALC_FLR.3:
 - first 5 as before
 - *The flaw remediation procedures shall include a procedure requiring timely response and the automatic distribution of security flaw reports and the associated corrections to registered users who might be affected by the security flaw.*
 - next 3 as before
 - *The flaw remediation guidance shall describe a means by which TOE users may register with the developer, to be eligible to receive security flaw reports and corrections.*
 - *The flaw remediation guidance shall identify the specific points of contact for all reports and enquiries about security issues involving the TOE.*

CEM -Common Evaluation Methodology

- given CC documentation, EAL classification etc, perform a check
- idea: evaluation by non-experts, semi-automated, mainly paper work
- mapping:
 - assurance class \Rightarrow activity
 - assurance component \Rightarrow sub-activity
 - evaluator action element \Rightarrow action
- responsibilities:
 - sponsor: requesting and supporting an evaluation. different agreements for the evaluation (e.g. commissioning the evaluation), providing evaluation evidence.
 - developer: produces TOE, providing the evidence required for the evaluation on behalf of the sponsor.
 - evaluator: performs the evaluation tasks required in the context of an evaluation, performs the evaluation sub-activities and provides the results of the evaluation assessment to the evaluation authority.
 - evaluation authority: establishes and maintains the scheme, monitors the evaluation conducted by the evaluator, issues certification/validation reports as well as certificates based on the evaluation results
- verdicts: pass, fail, inconclusive
- parts:
 - evaluation input task (are all documents available to perform evaluation?)
 - evaluation sub-activities
 - evaluation output task (deliver the Observation Report (OR) and the Evaluation Technical Report (ETR)).

- demonstration of the technical competence task

3 Security Problem Definition

- **Object Security Problem (OSP):** "The security problem definition defines the security problem that is to be addressed.
 - **axiomatic:** deriving the security problem definition outside the CC scope
 - **crucial:** the usefulness of the results of an evaluation strongly depends on the security problem definition.
 - **requires work:** spend significant resources and use well-defined processes and analyses to derive a good security problem definition.
- good example:

Secure signature-creation devices must, by appropriate technical and operational means, ensure at the least that:

 - 1) The signature-creation-data used for signature-creation can practically occur only once, and that their secrecy is reasonably assured;
 - 2) The signature-creation-data used for signature-creation cannot, with reasonable assurance, be derived and the signature is protected against forgery using currently available technology;
 - 3) The signature-creation-data used for signature-creation can be reliably protected by the legitimate signatory against the use of others
- **assets:** entities that someone places value upon. Examples of assets include: - contents of a file or a server; - the authenticity of votes cast in an election; - the availability of an electronic commerce process; - the ability to use an expensive printer; - access to a classified facility.

no threat no asset!
- **Threats:** threats to assets, what can happen that endangers assets
- **Assumptions:** assumptions are acceptable, where certain properties of the TOE environment are already known or can be assumed

this is NOT the place for putting properties derived from specific properties of the TOE

4. Security objectives

- "The security objectives are a concise and abstract statement of the intended solution to the problem defined by the security problem definition. Their role:
 - a high-level, natural language solution of the problem;
 - divide this solution into partwise solutions, each addressing a part of the problem;
 - demonstrate that these partwise solutions form a complete solution to the problem.
- bridge between the security problem and Security Functional Requirements (SFR)
- **mapping objectives to threats:** table, each threat should be covered, each objective has to respond to some threat

answers to questions:

- what is really needed?
- have we forgot about something?
- **rationale:** verifiable explanation why the mapping is sound

5. Extended Component Definition

- In many cases the security requirements (see the next section) in an ST are based on components in CC Part 2 or CC Part 3.
- in some cases, there may be requirements in an ST that are not based on components in CC Part 2 or CC Part 3.
- in this case new components (extended components) need to be defined

6.1 SFR (Security Functional requirements)

- *The SFRs are a translation of the security objectives for the TOE. They are usually at a more detailed level of abstraction, but they have to be a complete translation (the security objectives must be completely addressed) and be independent of any specific technical solution (implementation). The CC requires this translation into a standardised language for several reasons: - to provide an exact description of what is to be evaluated. As security objectives for the TOE are usually formulated in natural language, translation into a standardised language enforces a more exact description of the functionality of the TOE. - to allow comparison between two STs. As different ST authors may use different terminology in describing their security objectives, the standardised language enforces using the same terminology and concepts. This allows easy comparison.*
- predefined classes:
 - Logging and audit class FAU
 - Identification and authentication class FIA
 - Cryptographic operation class FCS
 - Access control families FDP_ACC, FDP_ACF
 - Information flow control families FDP_IFC, FDP_IFF
 - Management functions class FMT
 - (Technical) protection of user data families FDP_RIP, FDP_ITT, FDP_ROL
 - (Technical) protection of TSF data class FPT
 - Protection of (user) data during communication with external entities families FDP_ETC, FDP_ITC, FDP_UCT, FDP_UIT, FDP_DAU, classes FCO and FTP
- There is no translation required in the CC for the security objectives for the operational environment, because the operational environment is not evaluated
- customizing SFRs: **refinement** (more requirements), **selection** (options), **assignment** (values), **iterations** (the same component may appear at different places with different roles)
- rules:

check dependencies between SFR - In the CC Part 2 language, an SFR can have a dependency on other SFRs. This signifies that if an ST uses that SFR, it generally needs to use those other SFRs as well. This makes it much harder for the ST writer to overlook including necessary SFRs and thereby improves the completeness of the ST.

security objectives must follow from SFR's - Security Requirements Rationale section (Sect.6.3) in PP

if possible, use only standard SFR's

6.2 Security Assurance Requirements

- The SARs are a description of how the TOE is to be evaluated. This description uses a standardised language (to provide exact description, to allow comparison between two PP).
-

III. LEGAL FRAMEWORK - EXAMPLE: EIDAS REGULATION

goals:

- interoperability, comparable levels of trust
- merging national systems into pan-European one
- trust services, in particular: identification, authentication, signature, electronic seal, time-stamping, electronic delivery, Web authentication
- supervision system
- information about security breaches
- focused on public administration systems. However, the rules for all trust services except for closed systems (not available to anyone). Private sector encouraged to reuse the same means.

tools:

- common legal framework
- supervision system
- obligatory exchange of information about security problems
- common understanding of assurance levels

technical concept:

- each Member State provides an online system enabling identification and authentication with means from this Member State to be used abroad
- a notification scheme for national systems
- if notified (some formal and technical conditions must be fulfilled), then every member state must implement it in own country within 12 month

identification and authentication:

- eID cards – Member States are free to introduce any solution, the Regulation attempts to change it and build a common framework from a variety of (incompatible) solutions
- breakthrough claimed, but likely to fail

changes regarding electronic signature:

- electronic seal with the same conditions as electronic signature,
- the seal is aimed for legal persons
- weakening conditions for qualified electronic signatures: admitting server signatures and delegating usage of private keys

new:

- electronic registered delivery service
- Webpage authentication

Example of requirements (electronic seal):

Definition:

“electronic seal creation device” means configured software or hardware used to create an electronic seal;

“qualified electronic seal creation device” means an electronic seal creation device that meets mutatis mutandis the requirements laid down in Annex II;

Art. 36

An advanced electronic seal shall meet the following requirements:

- (a) it is uniquely linked to the creator of the seal;
- (b) it is capable of identifying the creator of the seal;
- (c) it is created using electronic seal creation data that the creator of the seal can, with a high level of confidence under its control, use for electronic seal creation; and
- (d) it is linked to the data to which it relates in such a way that any subsequent change in the data is detectable.

Annex II:

- (a) the confidentiality of the electronic signature creation data used for electronic signature creation is reasonably assured;
- (b) the electronic signature creation data used for electronic signature creation can practically occur only once;
- (c) the electronic signature creation data used for electronic signature creation cannot, with reasonable assurance, be derived and the electronic signature is reliably protected against forgery using currently available technology;
- (d) the electronic signature creation data used for electronic signature creation can be reliably protected by the legitimate signatory against use by others.

2. Qualified electronic signature creation devices shall not alter the data to be signed or prevent such data from being presented to the signatory prior to signing.

3. Generating or managing electronic signature creation data on behalf of the signatory may only be done by a qualified trust service provider.

4. Without prejudice to point (d) of point 1, qualified trust service providers managing electronic signature creation data on behalf of the signatory may duplicate the electronic signature creation data only for back-up purposes provided the following requirements are met:

(a) the security of the duplicated datasets must be at the same level as for the original datasets;

(b) the number of duplicated datasets shall not exceed the minimum needed to ensure continuity of the service.

Art. 30

1. Conformity of qualified electronic signature creation devices with the requirements laid down in Annex II shall be certified by appropriate public or private bodies designated by Member States.

notification system:

An electronic identification scheme eligible for notification if:

(a) issued by the notifying state

(b) at least one service available in this state;

(c) at least assurance level low;

(d) ensured that the person identification data is given to the right person

(e) ...

(f) availability of authentication online, for interaction with foreign systems (free of charge for public services), no specific disproportionate technical requirements

(g) description of that scheme published 6 months in advance

(h) meets the requirements from the implementing act

Assurance levels:

- regulation, Sept. 2015, implementation of eIDAS
- reliability and quality of
 - enrolment
 - electronic identification means management
 - authentication
 - management and organization
- authentication factors
 - possession based
 - knowledge based
 - inherent (physical properties)
- enrolment: (for all levels):
 1. Ensure the applicant is aware of the terms and conditions related to the use of the electronic identification means.

2. Ensure the applicant is aware of recommended security precautions related to the electronic identification means.
3. Collect the relevant identity data required for identity proofing and verification.

- identity proving and verification (for natural persons):

low:

1. The person can be assumed to be in possession of evidence recognised by the Member State in which the application for the electronic identity means is being made and representing the claimed identity.
2. The evidence can be assumed to be genuine, or to exist according to an authoritative source and the evidence appears to be valid.
3. It is known by an authoritative source that the claimed identity exists and it may be assumed that the person claiming the identity is one and the same.

substantial: low plus:

1. The person has been verified to be in possession of evidence recognised by the Member State in which the application for the electronic identity means is being made and representing the claimed identity

and

the evidence is checked to determine that it is genuine; or, according to an authoritative source, it is known to exist and relates to a real person

and

steps have been taken to minimise the risk that the person's identity is not the claimed identity, taking into account for instance the risk of lost, stolen, suspended, revoked or expired evidence; or

2. options related to other trustful sources

high: substantial plus

(a) Where the person has been verified to be in possession of photo or biometric identification evidence recognised by the Member State in which the application for the electronic identity means is being made and that evidence represents the claimed identity, the evidence is checked to determine that it is valid according to an authoritative source; and the applicant is identified as the claimed identity through comparison of one or more physical characteristic of the person with an authoritative source; or ...

- electronic identification means management:

low:

1. The electronic identification means utilises at least one authentication factor.
2. The electronic identification means is designed so that the issuer takes reasonable steps to check that it is used only under the control or possession of the person to whom it belongs.

substantial:

1. The electronic identification means utilises at least two authentication factors from different categories.
2. The electronic identification means is designed so that it can be assumed to be used only if under the control or possession of the person to whom it belongs.

high:

1. The electronic identification means protects against duplication and tampering as well as against attackers with high attack potential
2. The electronic identification means is designed so that it can be reliably protected by the person to whom it belongs against use by others.

- Issuance , delivery and activation:

low:

After issuance, the electronic identification means is delivered via a mechanism by which it can be assumed to reach only the intended person.

substantial:

After issuance, the electronic identification means is delivered via a mechanism by which it can be assumed that it is delivered only into the possession of the person to whom it belongs.

high:

The activation process verifies that the electronic identification means was delivered only into the possession of the person to whom it belongs.

- suspension, revocation and reactivation:

all levels:

1. It is possible to suspend and/or revoke an electronic identification means in a timely and effective manner.
2. The existence of measures taken to prevent unauthorised suspension, revocation and/or reactivation.
3. Reactivation shall take place only if the same assurance requirements as established before the suspension or revocation continue to be met.

- authentication mechanism:

substantial:

1. The release of person identification data is preceded by reliable verification of the electronic identification means and its validity.
2. Where person identification data is stored as part of the authentication mechanism, that information is secured in order to protect against loss and against compromise, including analysis offline.
3. The authentication mechanism implements security controls for the verification of the electronic identification means, so that it is highly unlikely that activities such as guessing, eavesdropping, replay or manipulation of communication by an attacker with enhanced-**basic attack potential** can subvert the authentication mechanisms.

high:

... by an attacker with **high attack potential** can subvert the authentication mechanisms.

- audit:

low:

The existence of periodical internal audits scoped to include all parts relevant to the supply of the provided services to ensure compliance with relevant policy.

substantial:

The existence of periodical **independent** internal or external audits

high:

1. The existence of periodical **independent external audits** scoped to include all parts relevant to the supply of the provided services to ensure compliance with relevant policy.
 2. Where a scheme is directly managed by a government body, it is audited in accordance with the national law.
-

III. TECHNICAL GOVERNMENT REGULATIONS - EXAMPLE :

FIPS PUB 140-2, SECURITY REQUIREMENTS FOR CRYPTOGRAPHIC MODULES

- Federal Information Processing Standards, NIST, recommendations and standards based on US law
- for sensitive but unclassified information
- levels: 1-4
- Cryptographic Module Validation Program (certification by NIST and Canadian authority)
- need to use “approved security functions” if to be used in public sector, waivers concerning some features are possible
- Levels:
 - Level 1: cryptographic module with at least one approved algorithm, no physical protection (like a PC)
 - Level 2:
 - tamper evident seals for access to CSP (critical security parameters)
 - role base authentication for operator,
 - refers to PPs, EAL2 or higheror secure operating system
 - Level 3:
 - protection against unauthorized access and attempts to modify cryptographic module, detection probability should be high,
 - CSP separated in a physical way from the rest
 - identity based authentication+ role based of an identified person (and not solely role based as on level 2)
 - CSP input and output - encrypted
 - components of cryptographic module can be executed in a general purpose operating system if
 - PP fulfilled, Trusted Path fulfilled

- EAL 3 or higher
 - security policy model (ADV.SPM1)
 - or a trusted operating system
- Level 4:
 - like level 3 but at least EAL4
- a more detailed overview:

	<i>Security Level 1</i>	<i>Security Level 2</i>	<i>Security Level 3</i>	<i>Security Level 4</i>
Cryptographic Module Specification	Specification of cryptographic module, cryptographic boundary, Approved algorithms, and Approved modes of operation. Description of cryptographic module, including all hardware, software, and firmware components. Statement of module security policy.			
Cryptographic Module Ports and Interfaces	Required and optional interfaces. Specification of all interfaces and of all input and output data paths.		Data ports for unprotected critical security parameters logically or physically separated from other data ports.	
Roles, Services, and Authentication	Logical separation of required and optional roles and services.	Role-based or identity-based operator authentication.	Identity-based operator authentication.	
Finite State Model	Specification of finite state model. Required states and optional states. State transition diagram and specification of state transitions.			
Physical Security	Production grade equipment.	Locks or tamper evidence.	Tamper detection and response for covers and doors.	Tamper detection and response envelope. EFP or EFT.
Operational Environment	Single operator. Executable code. Approved integrity technique.	Referenced PPs evaluated at EAL2 with specified discretionary access control mechanisms and auditing.	Referenced PPs plus trusted path evaluated at EAL3 plus security policy modeling.	Referenced PPs plus trusted path evaluated at EAL4.
Cryptographic Key Management	Key management mechanisms: random number and key generation, key establishment, key distribution, key entry/output, key storage, and key zeroization.			
	Secret and private keys established using manual methods may be entered or output in plaintext form.		Secret and private keys established using manual methods shall be entered or output encrypted or with split knowledge procedures.	
EMI/EMC	47 CFR FCC Part 15. Subpart B, Class A (Business use). Applicable FCC requirements (for radio).		47 CFR FCC Part 15. Subpart B, Class B (Home use).	
Self-Tests	Power-up tests: cryptographic algorithm tests, software/firmware integrity tests, critical functions tests. Conditional tests.			
Design Assurance	Configuration management (CM). Secure installation and generation. Design and policy correspondence. Guidance documents.	CM system. Secure distribution. Functional specification.	High-level language implementation.	Formal model. Detailed explanations (informal proofs). Preconditions and postconditions.
Mitigation of Other Attacks	Specification of mitigation of attacks for which no testable requirements are currently available.			

Table 1: *Summary of security requirements*

- more details:
 - roles: user, crypto officer, maintenance
 - services: to operator: show status, perform self-tests, perform approved security function, bypassing cryptographic operations must be documented etc.

- authentication: pbb of a random guess $< \frac{1}{1000000}$, one minute attempts: $< \frac{1}{100000}$, feedback obscured
- physical security:
 - full documentation,
 - if maintenance functionalities, then many features including erasing the key when accessed
 - protected holes, you cannot put probing devices through the holes
 - level 4: environmental failure protection (EFP) features or undergo environmental failure testing (EFT) – prevent leakage through unusual conditions

	General Requirements for all Embodiments	Single-Chip Cryptographic Modules	Multiple-Chip Embedded Cryptographic Modules	Multiple-Chip Standalone Cryptographic Modules
Security Level 1	Production-grade components (with standard passivation).	No additional requirements.	If applicable, production-grade enclosure or removable cover.	Production-grade enclosure.
Security Level 2	Evidence of tampering (e.g., cover, enclosure, or seal).	Opaque tamper-evident coating on chip or enclosure.	Opaque tamper-evident encapsulating material or enclosure with tamper-evident seals or pick-resistant locks for doors or removable covers.	Opaque enclosure with tamper-evident seals or pick-resistant locks for doors or removable covers.
Security Level 3	Automatic zeroization when accessing the maintenance access interface. Tamper response and zeroization circuitry. Protected vents.	Hard opaque tamper-evident coating on chip or strong removal-resistant and penetration resistant enclosure.	Hard opaque potting material encapsulation of multiple chip circuitry embodiment or applicable Multiple-Chip Standalone Security Level 3 requirements.	Hard opaque potting material encapsulation of multiple chip circuitry embodiment or strong enclosure with removal/penetration attempts causing serious damage.
Security Level 4	EFP or EFT for temperature and voltage.	Hard opaque removal-resistant coating on chip.	Tamper detection envelope with tamper response and zeroization circuitry.	Tamper detection/ response envelope with tamper response and zeroization circuitry.

Table 2: *Summary of physical security requirements*

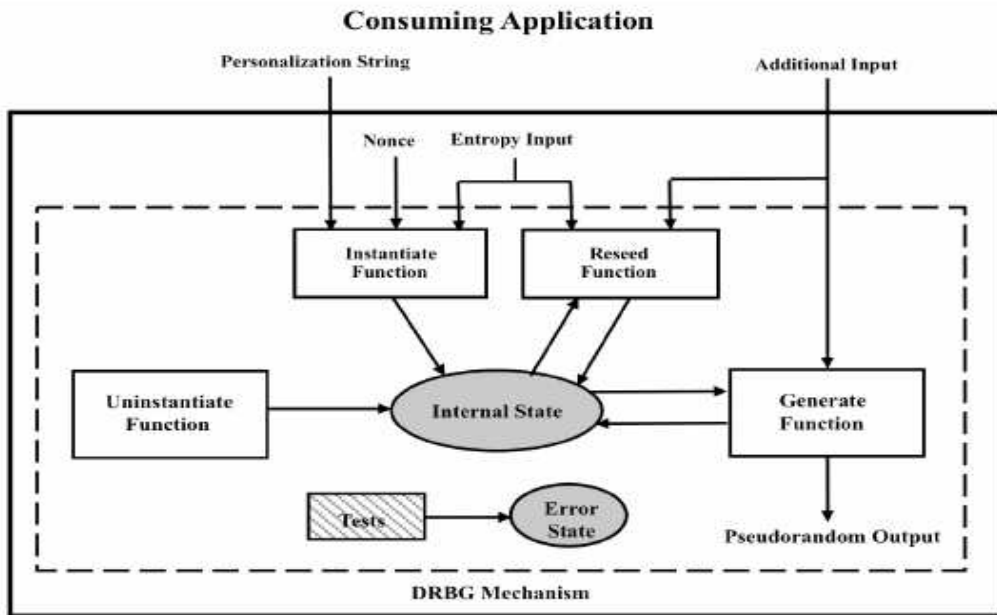
- more details:
 - operational environment:
 - L1: separation of processes, concurrent operators excluded, no interrupting cryptographic module, Approved integrity technique (HMAC?)
 - L2: operating system control functions under EAL2, specify roles to operate, modify,..., crypto software withing cryptographic boundary, audit: recording invalid operators, capable of auditing the following events:

- operations to process audit data from the audit trail,
 - requests to use authentication data management mechanisms,
 - use of a security-relevant crypto officer function,
 - requests to access user authentication data associated with the cryptographic module,
 - use of an authentication mechanism (e.g., login) associated with the cryptographic module,
 - explicit requests to assume a crypto officer role,
 - the allocation of a function to a crypto officer role.
 - L3: EAL3, trusted path (also included in audit trail)
 - L4: EAL4
 - key management:
 - non-approved RNG can be used for IV or as input to approved RNG
 - list of approved RNG: refers to an annex and annex to NIST document from 2016 (with a link to 2015)
 - list of approved key establishment - again links
 - key in out: automated (encrypted) or manual (splitted in L3 or L4)
 - tests: self-test and power-up. No crypto operation if something wrong. tests based on known outputs
 - Pair-wise consistency test (for public and private keys).
 - Software/firmware load test.
 - Manual key entry test.
 - Continuous random number generator test.
 - Bypass test – proper switching between bypass and crypto
-

FIPS Approved Random Number Generators

- nondeterministic generators not approved

- deterministic: special NIST Recommendation,
- first approved entropy source creates a seed , then deterministic part



Instantiation:

- seed has a limited period
- reseeded function requires a different seed
- different instantiations can exist at the same time, they MUST be independent in terms of the seeds and usage

Internal state:

- contains cryptographic chain value AND the number of requests so far
- different instantiations of DRBG must have separate internal states

Instantiation strength:

- formally defined as "112, 128, 192, 256 bits", intuition: number of bits to be guessed
- $\text{Security_strength_of_output} = \min(\text{output_length}, \text{DRBG_security_strength})$


functions:

- instantiate: initializing the internal state, preparing to use
- generate: generating output bits as PRNG
- reseed: combines the internal state with new entropy to change the seed

- uninstantiate: erase the internal state
- test: checks correctness of components on the chip

DRBG mechanism boundary:

- not cryptographic module boundary
- DRBG internal state and operation shall only be affected according to the DRBG mechanism specification
- the state exists solely within the DRBG mechanism boundary, not -accessible from outside
- information about internal state only via specified output



Seed:

- entropy is obligatory, entropy strength should be at least the entropy of the output
- reseeding: nonce not used, the internal state used
- *approved randomness source* obligatory for entropy source

- nonce: not secret. Example nonces:
 - a random value from an approved generator
 - a trusted timestamp of sufficient resolution (never use the same timestamp)
 - monotonically increasing sequence number
 - combination of a timestamp and a monotonically increasing sequence number, such that the sequence number is reset when and only when the timestamp change
- not used for any other purposes

reseed:

- "for security"
- argument: it might be better than `uninstantiate` and `instantiate` due to aging of randomness source

personalization:

- not security critical, but the adversary might be unaware of it (analogous to a login)

resistance:

- backtracking resistance: given internal state at time t it is infeasible to distinguish between the output for period $[1, t - 1]$ and a random output
- prediction resistance: *“Prediction resistance means that a compromise of the DRBG internal state has no effect on the security of future DRBG outputs. That is, an adversary who is given access to all of the output sequence after the compromise cannot distinguish it from random output with less work than is associated with the security strength of the instantiation; if the adversary knows only part of the future output sequence, he cannot predict any bit of that future output sequence that he does not already know (with better than a 50-50 chance). – refers only to reseeding (before reseeding the output is predictable) distinguishability from random input or predicting missing output bits*

specific functions:

- `(status, entropy_input) = Get_entropy_input (min_entropy, min_length, max_length, prediction_resistance_request)`,
- Instantiation:
 - checks validity of parameters
 - determines security strength
 - obtains entropy input, nonce
 - runs instantiate algorithm to get initial state
 - returns a handle

`Instantiate_function(requested_instantiation_security_strength,
prediction_resistance_flag, personalization_string)`

`prediction_resistance_flag` determines whether consuming application may request reseeding

- Reseed:

- Explicit request by a consuming application,
 - if prediction resistance is requested
 - also if the upper bound on the number of generated output reached
 - also due to external events
- steps:
- checks validity of the input parameters,
 - determines the security strength
 - obtains entropy input, nonce
 - runs reseed algorithm to get initial state

- Generate function:

`Generate_function(state_handle,requested_number_of_bits,requested_security_strength,
prediction_resistance_request, additional_input)`

- Removing a DRBG Instantiation:

`Uninstantiate_function (state_handle)`
internal state zeroized

Hash_DRBG

variants:

- hash algorithms: SHA-1 up to SHA-512
- parameters determined, e.g. maximum length of personalization string
- seed length typically 440 (but also 888)

state:

- value V updated during each call to the DRBG
- constant C that depends on the seed
- counter `reseed_counter`: storing the number of requests for pseudorandom bits since new entropy_input was obtained during instantiation or reseeding

instantiation:

1. `seed_material = entropy_input || nonce || personalization_string`
2. `seed = Hash_df (seed_material, seedlen)` (hash derivation function)
3. `V = seed`
4. `C = Hash_df ((0x00 || V), seedlen)`
5. Return (V, C, reseed_counter)

reseed:

1. `seed_material = 0x01 || V || entropy_input || additional_input`
2. `seed = Hash_df (seed_material, seedlen)`
3. `V = seed`
4. `C = Hash_df ((0x00 || V), seedlen)`
5. `reseed_counter = 1`
6. Return (V, C, reseed_counter).

generating bits:

1. If `reseed_counter > reseed_interval`, then return “reseed required”
2. If (`additional_input ≠ Null`), then do
 - 2.1 `w = Hash (0x02 || V || additional_input)`
 - 2.2 $V = (V + w) \bmod 2^{\text{seedlen}}$
3. `(returned_bits) = Hashgen (requested_number_of_bits, V)`
4. `H = Hash (0x03 || V)`
5. $V = (V + H + C + \text{reseed_counter}) \bmod 2^{\text{seedlen}}$
6. `reseed_counter = reseed_counter + 1`
7. Return (SUCCESS, returned_bits, V, C, reseed_counter)

Hashgen:

1. $m = \frac{\text{requested_no_of_bits}}{\text{outlen}}$
2. `data = V`
3. `W = Null string`
4. For `i = 1` to `m`

- 4.1 $w = \text{Hash}(\text{data})$.
- 4.2 $W = W \parallel w$
- 4.3 $\text{data} = (\text{data} + 1) \bmod 2^{\text{seedlen}}$
5. $\text{returned_bits} = \text{leftmost}(W, \text{requested_no_of_bits})$
6. Return (returned_bits).

HMAC_DRBG

Update (used for instantiation and reseeding)

1. $K = \text{HMAC}(K, V \parallel 0x00 \parallel \text{provided_data})$
2. $V = \text{HMAC}(K, V)$
3. If ($\text{provided_data} = \text{Null}$), then return K and V
4. $K = \text{HMAC}(K, V \parallel 0x01 \parallel \text{provided_data})$
5. $V = \text{HMAC}(K, V)$
6. Return (K, V).

Instantiate:

1. $\text{seed_material} = \text{entropy_input} \parallel \text{nonce} \parallel \text{personalization_string}$
2. $\text{Key} = 0x00\ 00\dots00$
3. $V = 0x01\ 01\dots01$
4. $(\text{Key}, V) = \text{HMAC_DRBG_Update}(\text{seed_material}, \text{Key}, V)$
5. $\text{reseed_counter} = 1$
6. Return ($V, \text{Key}, \text{reseed_counter}$)

Reseed:

1. $\text{seed_material} = \text{entropy_input} \parallel \text{additional_input}$
2. $(\text{Key}, V) = \text{HMAC_DRBG_Update}(\text{seed_material}, \text{Key}, V)$
3. $\text{reseed_counter} = 1$
4. Return ($V, \text{Key}, \text{reseed_counter}$).

Generate bits:

1. If $\text{reseed_counter} > \text{reseed_interval}$, then return “reseed required”
2. If $\text{additional_input} \neq \text{Null}$, then

- (Key, V) = HMAC_DRBG_Update (additional_input, Key, V)
3. temp = Null
 4. While len (temp) < requested_number_of_bits do:
 - 4.1 V = HMAC (Key, V)
 - 4.2 temp = temp || V
 5. returned_bits = leftmost (temp, requested_number_of_bits)
 6. (Key, V) = HMAC_DRBG_Update (additional_input, Key, V)
 7. reseed_counter = reseed_counter + 1
 8. Return (SUCCESS, returned_bits, Key, V, reseed_counter).

CTR_DRBG

a generator based on an encryption function, AES versions

internal state:

- value V of blocklen bits, updated each time another blocklen bits of output are produced
- keylen-bit Key, updated whenever a predetermined number of output blocks are generated
- counter (reseed_counter) = the number of requests for pseudorandom bits since instantiation or reseeding
- ctr_len is a parameter known by implementation

Update Process:

1. temp = Null
2. While (len (temp) < seedlen, do
 - 2.1 If ctr_len < blocklen
 - 2.1.1 inc = (rightmost (V, ctr_len) + 1) mod $2^{\text{ctr_len}}$.
 - 2.1.2 V = leftmost (V, blocklen-ctr_len) || inc
 - Else V = (V+1) mod 2^{blocklen}
 - 2.2 output_block = Block_Encrypt (Key, V)
 - 2.3 temp = temp || output_block
3. temp = leftmost (temp, seedlen)
4. temp = temp \oplus provided_data
5. Key = leftmost (temp, keylen)
6. V = rightmost (temp, blocklen).

Instantiate:

1. pad `personalization_string` with zeroes
2. `seed_material = entropy_input \oplus personalization_string`
3. `Key = 0keylen`
4. `V = 0blocklen`
5. `(Key, V) = CTR_DRBG_Update (seed_material, Key, V).`
6. `reseed_counter = 1`
7. Return `(V, Key, reseed_counter)`.

reseeding is similar

Generate:

1. If `reseed_counter > reseed_interval`, then “reseed required”
2. If `(additional_input \neq Null)`, then
 - 2.1 `temp = len (additional_input).`
 - 2.2 If `(temp < seedlen)` then pad `additional_input` with zeroes
 - 2.3 `(Key, V) = CTR_DRBG_Update (additional_input, Key, V).`
Else `additional_input = 0seedlen`
3. `temp = Null`
4. While `(len (temp) < requested_number_of_bit)`, do
 - 4.1 If `ctr_len < blocklen`
 - 4.1.1 `inc = (rightmost (V, ctr_len) + 1) mod 2ctr_len`
 - 4.1.2 `V = leftmost (V, blocklen-ctr_len) || inc`
 - Else `V = (V+1) mod 2blocklen`
 - 4.2 `output_block = Block_Encrypt (Key, V).`
 - 4.3 `temp = temp || output_block`
5. `returned_bits = leftmost (temp, requested_number_of_bits)`
6. `(Key, V) = CTR_DRBG_Update (additional_input, Key, V)`
7. `reseed_counter = reseed_counter + 1`
8. Return `(SUCCESS, returned_bits, Key, V, reseed_counter)`.

IV. GDPR

Some discussion from the lecture missing here

DOMAIN SIGNATURES

Definition of Domain Signatures

some discussion from the lecture missing here

Domain signatures and applications

1. replacing login+ password mechanisms. Like Restricted Identification from German Personal Identity cards. Advantages
 1. one PIN per person
 2. high entropy, not predictable
 3. no problem for
2. anonymization in datasets
3. recommendations,
4. "randomness" and self-organization

Pseudonymous Signature:

- BSI standard
- keys:
 - domain parameters D_M and a pair of global keys (PK_M, SK_M)
 - public key PK_{ICC} for a group of eIDAS tokens, the private key SK_{ICC} known to the issuer of eIDAS tokens (called manager)
 - for a token the manager chooses $SK_{ICC,2}$ at random, then computes $SK_{ICC,1}$ such that $SK_{ICC} = SK_{ICC,1} + SK_M \cdot SK_{ICC,2}$
 - a sector (domain) holds private key SK_{sector} and public key PK_{sector} .

- a sector has revocation private key $\text{SK}_{\text{revocation}}$ and public key $\text{PK}_{\text{revocation}}$
- sector specific identifiers $I_{\text{ICC},1}^{\text{sector}}$ and $I_{\text{ICC},2}^{\text{sector}}$ of the eIDAS token in the sector:
 $I_{\text{ICC},1}^{\text{sector}} = (\text{PK}_{\text{sector}})^{\text{SK}_{\text{ICC},1}}$ and $I_{\text{ICC},2}^{\text{sector}} = (\text{PK}_{\text{sector}})^{\text{SK}_{\text{ICC},2}}$
- signing: with keys $\text{SK}_{\text{ICC},1}$, $\text{SK}_{\text{ICC},2}$ and $I_{\text{ICC},1}^{\text{sector}}$ and $I_{\text{ICC},2}^{\text{sector}}$ for $\text{PK}_{\text{sector}}$ and message m
 - i. choose K_1, K_2 at random
 - ii. compute
 - $Q_1 = g^{K_1} \cdot (\text{PK}_M)^{K_2}$
 - $A_1 = (\text{PK}_{\text{sector}})^{K_1}$
 - $A_2 = (\text{PK}_{\text{sector}})^{K_2}$
 - iii. $c = \text{Hash}(Q_1, I_{\text{ICC},1}^{\text{sector}}, A_1, I_{\text{ICC},2}^{\text{sector}}, A_2, \text{PK}_{\text{sector}}, m)$
(variant parameters and $\mathbf{\Pi}$ omitted here)
 - iv. compute
 - $s_1 = K_1 - c \cdot \text{SK}_{\text{ICC},1}$
 - $s_2 = K_2 - c \cdot \text{SK}_{\text{ICC},2}$
 - v. output (c, s_1, s_2)
- verification:

compute

 - $Q_1 = (\text{PK}_{\text{ICC}})^c \cdot g^{s_1} \cdot (\text{PK}_M)^{s_2}$
 - $A_1 = (I_{\text{ICC},1}^{\text{sector}})^c \cdot (\text{PK}_{\text{sector}})^{s_1}$
 - $A_2 = (I_{\text{ICC},2}^{\text{sector}})^c \cdot (\text{PK}_{\text{sector}})^{s_2}$
 - recompute c and check against the c from the signature
- why it works?
$$\begin{aligned} (\text{PK}_{\text{ICC}})^c \cdot g^{s_1} \cdot (\text{PK}_M)^{s_2} &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c \cdot \text{SK}_{\text{ICC},1}} \cdot (\text{PK}_M)^{c \cdot \text{SK}_{\text{ICC},2}} \\ &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c \cdot \text{SK}_{\text{ICC},1}} \cdot (g)^{-c \cdot \text{SK}_M \cdot \text{SK}_{\text{ICC},2}} \\ &= (\text{PK}_{\text{ICC}})^c \cdot g^{K_1} \cdot (\text{PK}_M)^{K_2} \cdot g^{-c \cdot \text{SK}_{\text{ICC}}} = g^{K_1} \cdot (\text{PK}_M)^{K_2} = Q_1 \end{aligned}$$
- there is a version without A_1, A_2 and the pseudonyms $I_{\text{ICC},1}^{\text{sector}}, I_{\text{ICC},2}^{\text{sector}}$
- Problems:
 - the authorities know the private keys (there is a way to solve it when the user gets two pairs of keys on the device and takes their linear combination)
 - breaking into just 2 devices reveals the system keys

V. STANDARDS VERSUS SECURITY

It is not true that a standard solution is by definition a secure solution.

Standardisation process:

- representatives of countries, not necessarily specialists
- represent interests of industry
- target: a unified solution
- no open evaluation as in case of NIST competitions
- long process, many standards never used in practice

result: no guarantee for security

Example: Bleichenbacher's RSA signature forgery based on implementation error

- background: one has to pad the hash value before applying RSA signing operation, after padding the input no more a small number
- PKCS#1: RSA Laboratories standard for formats of RSA signatures (currently 1.5 and OEAP (optimal Asymmetric encryption padding))
- The attack works for PKCS-1.5 padding:
 - a byte 0
 - a byte 1
 - string of 0xFF bytes (their number depends on RSA modulus and the rest)
 - a byte 0
 - ASN.1 data
 - hash

```
00 01 FF FF FF ... FF 00 ASN.1 HASH
```
- RSA signature verification: exponentiation with the public key, remove padding, check the hash

- implementation based on the standard: recognize the structure 00 01 FF FF FF ... FF 00 and after them parse the hash
- attack mechanism:
 - hash not right adjusted (padding short), after the hash there is a part that is not parsed (it could be anything)
 - concern RSA systems with public key 3 (sometimes it is done so to speed-up verification) – according to strong RSA assumption computing roots of degree 3 is generally infeasible without the secret key
 - part after the hash adjusted so that the resulting number is a cube as an integer
 - compute the root ... and this is the signature!
- attack variants: some fields declared but not checked. then Bleichenbacher's freedom to adjust the number to become a cube even if the hash is right justified

Chosen Ciphertext Attacks Against Protocols Based on RSA Encryption Standard PKCS-1 – the Million Message Attack.

- RSA decryption device, returns an error message if the ciphertext not in PKCS-1 format (HSM,...)
- the ciphertext c to be broken is manipulated many times and based on error messages we narrow the set of choices for the plaintext
- attack (find m such that $m^d = c \pmod n$):
 1. phase: blind the ciphertext: $c_0 := c \cdot s^e \pmod n$ by choosing s such that c_0 is a valid PKCS-1 ciphertext.
 - Observe $c_0^e = m \cdot s$ and it starts with msb: 00, 02,
 - let k the byte length of n , $B = 2^{8(k-2)}$
 - then $2B \leq m \cdot s < 3B$
 - let $M_0 = [2B, 3B - 1]$
 2. phase: narrowing the set of intervals defining $s_1 < s_2 < \dots$ and M_1, M_2, \dots such that $M_{i+1} \subset M_i$, each M_i is a set of intervals
 - if M_{i-1} is a single interval $[a, b]$ then choose small values s_i and r_i such that $\frac{2B + r_i n}{b} < s_i < \frac{3B + r_i n}{a}$ and $c_0 \cdot s_i^e$ is PKCS-1 conforming
 - if M_{i-1} is not a single interval, then simply smallest s_i such that $c_0 \cdot s_i^e$ is PKCS-1 conforming
 - M_i consists of all intervals $[\max(a, \frac{2B + r_i \cdot n}{s_i}), \min(b, \frac{3B - 1 + r_i \cdot n}{s_i})]$ for $[a, b]$ from M_{i-1} and

$$\frac{a \cdot s_i - 3B + 1}{n} \leq r \leq \frac{b \cdot s_i - 2B}{n}$$

- when eventually $M_i = [a, a]$ then set $m = a \cdot s_0^{-1} \bmod n$

The attack exploits such vulnerabilities like

- error messages returned by the attacked device when decryption fails on different stages of the decryption algorithm
- different timings of execution of the decryption algorithm when the PKCS-1 encryption padding is correct and when it is incorrect.

If a device supports the PKCS-1 encryption padding and the implementation of the PKCS-1 decryption on the device is vulnerable, then the million message attack works also when

- the ciphertext is calculated according to a padding different than PKCS-1
- the “ciphertext” is the plaintext for which we want to obtain a signature (dangerous for a situation when the same key is used for decryption and for signatures, and decryption is not PIN protected).

VI. RFC DOCUMENTS (actual lecture given by dr Kubiak could be different)

RFC

”Request for Comments”

- by Internet Engineering Task Force (IETF) and the Internet Society
- semi-standard, developed from rfc from ARPANET
- authors of RFC versus standards with committees
- peer review, some reach status of “Internet Standards”
- RFC editor provided
- streams:
 - Internet Engineering Task Force (IETF) - current issues
 - BCP Best Current Practice;
 - FYI For Your Information; informational
 - STD Standard: with 2 maturity levels
 - Internet Research Task Force (IRTF) - more long term issues
 - Internet Architecture Board (IAB) (a body over task forces)

- independent
- Status:
 - informational
 - experimental
 - best current practice
 - standard: Proposed Standard, Draft Standard, Internet Standard

RFC as an example of specification of a protocol

A. Exemplary RFC: draft of TLS 1.3 spec.

B. Required level of detail - ensuring unambiguous implementation.

C. Structure of the document: from general view to detailed description - to facilitate reading many datastructures and technicalities are shifted to the appendices.

- Abstract: What the document is about?
- Status: Internet draft, expiration date
- Copyright notice
- Table of contents:

1. Introduction

- a very nice note for RFC editors (present in the draft only)
- the goal of the protocol (authentication, confidentiality, integrity + definitions of the three terms, to let the non-cryptographers understand the document)
- the high level view - primary components
- conventions and terminology - for clear and precise understanding
- change log - for editors (present in the draft only)
- major differences from previous version of the protocol (TLS 1.2)

2. Protocol Overview. Handshake: what must be negotiated, what are the basic key exchange modes?

3. Presentation language: Big or little endian? basic block size? etc.

4.-9. Protocol components, further details.

10. Security considerations

D. Exemplary protocol detail: certificate request from server side (dnames of CAs) cross-certification.

EXAMPLE: RFC2560

Network Working Group

Category: Standards Track

authors ... June 1999

title: *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies a protocol useful in determining the current status of a digital certificate without requiring CRLs. Additional mechanisms addressing PKIX operational requirements are specified in separate documents.

... contents of sections

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document (in uppercase, as shown) are to be interpreted as described in [RFC2119].

—

MUST=REQUIRED=SHALL: an absolute requirement

MUST NOT=SHALL NOT: an absolute prohibition of the specification

SHOULD=RECOMMENDED: ‘*there may exist valid reasons in particular circumstances to ignore, but implications must be understood and carefully weighed before choosing a different course*’

SHOULD NOT=NOT RECOMMENDED: negation of SHOULD (think twice before implementing it in this way!)

MAY=OPTIONAL: real option, **but** implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option,

2. Protocol Overview

- supplement to periodical checking CRL
- enables to determine the state of an identified certificate
- more timely, with more information
- RFC defines data exchanged

2.1 Request

– protocol version – service request – target certificate identifier – optional extensions which MAY be processed by the OCSP Responder

OCSP Responder checks:

1. request well formed
2. responder configured to serve such request
3. all necessary data given in the request

otherwise: error message

2.2 Response

- type+actual response
- basic type MUST be supported
- *"All definitive response messages SHALL be digitally signed."*
- signer MUST be one of: CA who issued the certificate, or a Trusted Responder of the requester, CA Designated Responder (Authorized Responder) - agent of CA with a certificate from CA
- response message: version of the response syntax – name of the responder – responses for each of the certificates in a request – optional extensions – signature algorithm OID – signature computed across hash of the response
- for each target certificate: certificate status value – response validity interval – optional extensions
- values:
 - good: *"At a minimum, this positive response indicates that the certificate is not revoked, but does not necessarily mean that the certificate was ever issued or that the time at which the response was produced is within the certificate's validity interval. Response extensions may be used to convey additional information on assertions made by the responder regarding the status of the certificate such as positive statement about issuance, validity, etc."*
 - revoked: the certificate has been revoked (permanently or temporarily (on hold))
 - unknown: responder has no data

2.3 Exception Cases

- error messages not signed
- types: – malformedRequest – internalError – tryLater – sigRequired – unauthorized
- "internalError" = responder reached an inconsistent internal state. The query should be retried
- "tryLater" = temporarily unable to respond

- "sigRequired"= the server requires the client sign
- "unauthorized"=the client is not authorized to make this query

2.4 Semantics of thisUpdate, nextUpdate and producedAt

- thisUpdate = time at which the indicated status is known to be correct
- nextUpdate= time at or before which newer information will be available about the certificate status
- producedAt = time at which the OCSP signed this response.

2.5 Response Pre-production

“OCSP responders MAY pre-produce signed responses specifying the status of certificates at a specified time. The time at which the status was known to be correct SHALL be reflected in the thisUpdate field of the response. The time at or before which newer information will be available is reflected in the nextUpdate field, while the time at which the response was produced will appear in the producedAt field of the response.”

- means that OCSP is not checking the status of the certificate but status on the CRL!

2.6 OCSP Signature Authority Delegation

- the OCSP might be an agent of CA explicitly appointed,
- signing key must allow signing it

2.7 CA Key Compromise

- if CA's private key compromised, then OCSP MAY return the revoked state for all certificates issued by that CA.

3. Functional Requirements

3.1 Certificate Content

- CAs SHALL provide the capability to include the AuthorityInfoAccess extension in certificates that can be checked using OCSP
- accessLocation for the OCSP provider may be configured locally at the OCSP client
- CAs supporting OCSP MUST “provide for the inclusion of a value for a uniformResourceIndicator (URI) accessLocation and the OID value id-ad-ocsp for the accessMethod in the AccessDescription SEQUENCE”
- accessLocation field in the subject certificate defines the transport (e.g. HTTP) used to access OCSP responder and data (e.g. a URL)

3.2 Signed Response Acceptance Requirements

Before accepting response clients SHALL confirm that:

1. certificate in response=certificate asked
2. signature valid

3. signature of the responder
4. responder authorized
5. thisUpdate sufficiently recent
6. nextUpdate is greater than the current time

4. Detailed Protocol

- data to be signed encoded using ASN.1 distinguished encoding rules (DER)
- ASN.1 EXPLICIT tagging as a default
- ”terms imported from elsewhere are: Extensions, CertificateSerialNumber, SubjectPublicKeyInfo, Name, AlgorithmIdentifier, CRLReason”

4.1 Requests

4.1.1 Request Syntax

OCSPRequest ::= SEQUENCE { tbsRequest TBSRequest, optionalSignature [0] EXPLICIT Signature OPTIONAL }

TBSRequest ::= SEQUENCE { version [0] EXPLICIT Version DEFAULT v1, requestorName [1] EXPLICIT GeneralName OPTIONAL, requestList SEQUENCE OF Request, requestExtensions [2] EXPLICIT Extensions OPTIONAL }

Signature ::= SEQUENCE { signatureAlgorithm AlgorithmIdentifier, signature BIT STRING, certs [0] EXPLICIT SEQUENCE OF Certificate OPTIONAL }

Version ::= INTEGER { v1(0) }

Request ::= SEQUENCE { reqCert CertID, singleRequestExtensions [0] EXPLICIT Extensions OPTIONAL }

CertID ::= SEQUENCE { hashAlgorithm AlgorithmIdentifier, issuerNameHash OCTET STRING, – Hash of Issuer’s DN issuerKeyHash OCTET STRING, – Hash of Issuers public key serialNumber CertificateSerialNumber }

- public key hashed together with name (names may repeat, public key must not)
- Support for any specific extension is OPTIONAL
- ”Unrecognized extensions MUST be ignored (unless they have the critical flag set and are not understood)”.
- requestor MAY sign the OCSP request, data included for easy verification (name:SHALL, certificate: MAY)

4.2 Response Syntax

OCSPResponse ::= SEQUENCE { responseStatus OCSPResponseStatus, responseBytes [0] EXPLICIT ResponseBytes OPTIONAL }

OCSPResponseStatus ::= ENUMERATED { successful (0), –Response has valid confirmations malformedRequest (1), –Illegal confirmation request internalError (2), –Internal error in issuer tryLater (3), –Try again later –(4) is not used sigRequired (5), –Must sign the request unauthorized (6) –Request unauthorized }

The value for `responseBytes` consists of an OBJECT IDENTIFIER and a response syntax identified by that OID encoded as an OCTET STRING.

`ResponseBytes ::= SEQUENCE { responseType OBJECT IDENTIFIER, response OCTET STRING }`

For a basic OCSP responder, `responseType` will be `id-pkix-ocsp-basic`.

`id-pkix-ocsp OBJECT IDENTIFIER ::= { id-ad-ocsp } id-pkix-ocsp-basic OBJECT IDENTIFIER ::= { id-pkix-ocsp 1 }`

4.3 Mandatory and Optional Cryptographic Algorithms

- clients SHALL: DSA sig-alg-oid specified in section 7.2.2 of [RFC2459]
- clients SHOULD: RSA signatures as specified in section 7.2.1 of [RFC2459]
- responders SHALL: SHA1

4.4 Extensions

4.4.1 Nonce

nonce against replay:

- nonce as one of the requestExtensions in requests
- in responses it would be included as one of the responseExtensions
- object identifier `id-pkix-ocsp-nonce`

4.4.2 CRL References

if revoked then indicate CRL where revoked

`id-pkix-ocsp-crl OBJECT IDENTIFIER ::= { id-pkix-ocsp 3 }`

`CrlID ::= SEQUENCE { crlUrl [0] EXPLICIT IA5String OPTIONAL, crlNum [1] EXPLICIT INTEGER OPTIONAL, crlTime [2] EXPLICIT GeneralizedTime OPTIONAL }`

For the choice `crlUrl`, the `IA5String` will specify the URL at which the CRL is available. For `crlNum`, the `INTEGER` will specify the value of the CRL number extension of the relevant CRL. For `crlTime`, the `GeneralizedTime` will indicate the time at which the relevant CRL was issued.

4.4.3 Acceptable Response Types

`d-pkix-ocsp-response OBJECT IDENTIFIER ::= { id-pkix-ocsp 4 }`

`AcceptableResponses ::= SEQUENCE OF OBJECT IDENTIFIER`

4.4.4 Archive Cutoff

- specifies how many years after expiration the revocation information is retained, this is "archive cutoff" date

4.4.5 CRL Entry Extensions

All the extensions specified as CRL Entry Extensions - in Section 5.3 of [RFC2459] - are also supported as `singleExtensions`.

4.4.6 Service Locator

OCSP server receives a request and reroutes it to another OCSP

serviceLocator request extension used

d-pkix-ocsp-service-locator OBJECT IDENTIFIER ::= { id-pkix-ocsp 7 }

ServiceLocator ::= SEQUENCE { issuer Name, locator AuthorityInfoAccessSyntax OPTIONAL }

Values defined in certificate asked

5. Security Considerations

- flood of queries,
- signed and unsigned both enable DOS
- precomputation helps
- HTTP caching might be risky: *“Implementors are advised to take the reliability of HTTP cache mechanisms into account when deploying OCSP over HTTP.”*

VII. HARDWARE TROJANS

methods of testing:

- functional tests
- internal tests circuitry
- optical inspection (destructive) - can detect modifications on layout level

Idea: change properties that are not visible under microscope: increase aging effects, manipulate transistors so that the output is fixed

Dopant Trojans

CMOS inverter: (image Wikipedia)




Figure 1.

where: A is the source, Vdd positive supply , Vss is ground

upper transistor: PMOS (allows current flow at low voltage)

lower transistor: NMOS (allows current flow at high voltage)


how it works:

- if voltage is low then the lower transistor (NMOS is in high resistance state and the current from Q flows to Vdd (high voltage)
- if voltage is high then the upper transistor (PMOS) is in high resistance state and the current from Q flows to Vss while Vdd has low voltage


PMOS: in dopant area “holes” (positive) playing the role of conductor, low voltage creates depletion area, high voltage attracts them

NMOS: in dopant area electrons (negative) playing the role of conductor, high voltage pushes the electrons out

CMOS inverter in the “bird eye perspective”:



Trojan design:



- whatever happens the VDD is connected to the output

Trojan TRNG

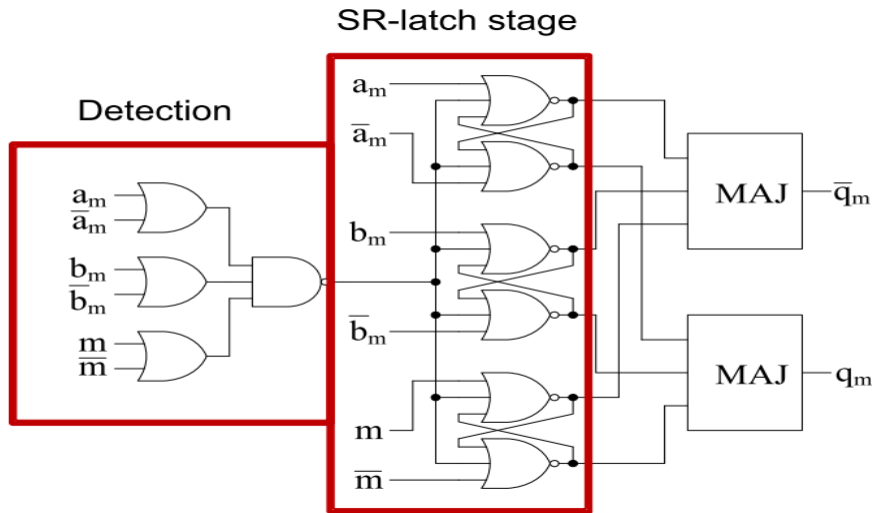
TRNG consists of

- entropy source (physical)
- self test circuit (OHT - online health test)
- deterministic RNG, Intel version:
 - conditioner (computes seeds to rate matcher) and rate matcher (computes 128 bit numbers)
 - derivation, internal state (K, c) :
 1. $c := c + 1, r := \text{AES}_K(c)$
 2. $c := c + 1, x := \text{AES}_K(c)$
 3. $c := c + 1, y := \text{AES}_K(c)$
 4. $K := K \oplus x$
 5. $c := c \oplus y$
- attack: fix K by applying Trojan transistors, if K is known, then it is easy to find internal state c from r and then the consecutive random numbers r

- problem with OHT: tests with some values have to create known outputs (32 CRC from the last 4 outputs), knowing the test one can find K by exhaustive search

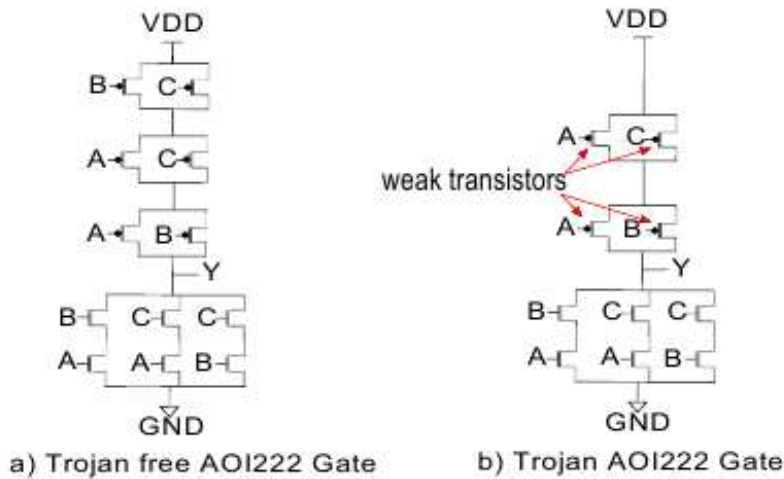
Side channel Trojan:

- side channel resistant logic: Masked Dual Rail Logic
 - for each a both a and negation of a computed
 - precharge: each phase preceded by charging all gates
 - masking operations by random numbers:
 - computing $a \wedge b$:
 - input $a \oplus m, a \oplus \neg m, b \oplus m, b \oplus \neg m, m, \neg m$
 - detection, SR-latch stage and majority gate



gates on the picture: OR, AND, NOR (output 1 if all inputs 0)

attacking not-majority gate:



Idea: instead of cutting output a low voltage

- the same behavior except for $A = 0$ and $B, C = 1$, where good output but high power consumption due to connection between VDD and VSS

Artificial aging

make some transistors disfunctional (as ithe case of PRNG)

method:

- apply to high voltage at certain areas
- the electrons accelerate and break barrier - damages
- effect the same as of aging a chip
- the transistor changes its operational characteristic

Defense methods:

- problem: Trojan may be triggered by some particular event, detection becomes harder
- problem: Trojan may work in very particular physical conditions, e.g. temperature, voltage
- on-chip checks: detection of unexpected behavior, e.g. delay characteristics: workload path and a shadow path that provides result after fixed time, + comparison
- ring osscillators on the chip detecting nonstandard behavior
- methods to enable activation in certain areas only
- inserting PUFs, (either randomize as much a s possible - noise over trojan information)
- keep algorithm deterministic
- secure coding: take into account the situation that certain components are not working properly
- external watchdogs techniques

VIII. QUANTUM COMPUTING AND QUANTUM DEVICES

as there are problems to guarantee security of devices in the traditional way, maybe there is a way out using physics? three directions:

- 1) quantum based random number generators
- 2) key transport
- 3) quantum cryptanalysis

Qubit

the concept is as follows:

- instead of a bit with discrete states 0 and 1 we have a linear combination of basis vectors denoted by $|0\rangle$ and $|1\rangle$:

$$\alpha \cdot |0\rangle + \beta \cdot |1\rangle$$
- with α, β complex numbers
- a measurement of $\alpha \cdot |0\rangle + \beta \cdot |1\rangle$ yields $|0\rangle$ with pbb $|\alpha|^2$ and $|1\rangle$ with pbb $|\beta|^2$ – this is quite annoying but ...
- measurement may be performed only for orthogonal basis. the basis can be different from $|0\rangle$ and $|1\rangle$. E.g.:

$$\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad , \quad \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$$

- measuring $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ in basis $|0\rangle$ and $|1\rangle$ yields both 0 and 1 with perfect probability 0.5: it seems to be a perfect source of random bits:
 - generate fotons $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$
 - measure them in basis $|0\rangle$ and $|1\rangle$
- moreover: reading changes the state to the state read: if the result is $|0\rangle$ then the physical state becomes $|0\rangle$ as well. There is no state $\alpha \cdot |0\rangle + \beta \cdot |1\rangle$ anymore.
- **In fact, this is the core of Shor’s algorithm - a reading operation creates a change in a physical system that would be infeasible to compute on a classical computer**
- instead of a single bit we may have strings of qubits, say of length l where $l > n$

Random Number Generators

Problems:

- high price (1000 EUR and more)
- while physical source might be ok, reading circuit introduces high bias, very poor results (2017) in the standard randomness tests for devices available on the market
- bias can be removed via additional logic, but extra hardware may mean place for Trojans and the whole advantage is gone

Quantum key transport, BB84

- Charles Bennett and Gilles Brassard, 1984, 1st quantum protocol, even implemented
- key agreement immune to eavsdropping (reading qubits is detectable)
- two bases used: $|0\rangle$ and $|1\rangle$ (denoted +) or $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ (denoted \times)
- encoding of bits:
 basis +: $|0\rangle = 0$ and $|1\rangle = 1$

basis \times : $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = 0, \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = 1$

– Steps:

1. Alice chooses at random bitstrings a and b of length n
2. for $i \leq n$ Alice encodes a_i according to basis indicated by b_i (0 indicates $+$, 1 indicates \times)
3. Alice sends n photons (codes for a)
4. Bob chooses at random string \hat{b} of n bits,
5. Bob measures the photons using basis indicated by \hat{b}_i for the i th photon
6. Alice sends b to Bob over a traditional (public) channel, Bob sends \hat{b} to Alice
7. Alice and Bob take the substring K of bits a_i such that $b_i = \hat{b}_i$
8. Alice chooses a subset of 50% of bit of K and discloses them to Bob
9. Bob checks how many of them disagree with his measurement. If above some threshold then it is likely that an adversary has measured the transmission as well and the protocol is aborted (environment may also create inconsistencies)
10. the unpublished substring of K may differ between Alice and Bob: an error correcting procedure applied (error correction attracts the bitstrings to the closest codewords, so if the strings of Alice and Bob differ slightly, then they result in the same codeword)
11. "privacy amplification": hashing to a much shorter string

– effect of eavesdropping:

- say Alice chooses basis \times to encode 0,
- eavesdropper Eve chooses a different basis for the measurement: namely $+$
- Eve gets $|0\rangle$ with pbb .5 and $|1\rangle$ with pbb .5, say $|0\rangle$ has been obtained
- at the same time the photon converted to $|0\rangle$!!!!!
- Bob measures the photon according to basis $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$
- $|0\rangle = \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) + \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)\right)$, so both results of the measurements (i.e. $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ or $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$) are equally probably for Bob, so the measurement of Bob indicated 1 with pbb .5
- corollary: eavesdropping creates inconsistency between Alice and Bob with pbb .5 once Eve chooses a different basis than Alice and Bob

– quantum hacking: in theory wonderful, but the problems come with physical realization

- sending many photons to Bob at the time when his hardware already set for a measurement. reflected photons show the basis used

Eckert algorithm:

- entangled pair of photons: measurement of one of them makes the mirror change of the other photon
- procedure:
 1. generate entangled qubits $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ by some source
 2. measure them by Alice and Bob on both ends of the channel

Properties:

- long distance transmissions possible, even to moving airplane
- over optical fibre or in free space (better vacuum)
- 1203 km between ground stations over satellite (China)
- both BB84 and Eckert used
- high price
- does not solve man-in-the-middle issue

Quantum computing and Shor factorization algorithm

Problem and its algebraic context:

- given an RSA number $n = p \cdot q$ for prime factors p and q of a similar size, the goal is to find p or q
- many modern crypto products are based on difficulty of this factorization problem. There are many software systems and embedded devices with RSA, no update is possible
- in order to break factorization problem it suffices to learn a nontrivial root r of 1:
 - $r \neq -1$
 - $r^2 = 1 \pmod n$

indeed

- $r^2 - 1 = (r - 1)(r + 1) = 0 \pmod{p \cdot q}$
- therefore p divides either $r - 1$ or $r + 1$
- if p divides $r - 1$ then q cannot divide $r - 1$ as then $r - 1$ would be at least n , but $r - 1 < n$
- in this situation we compute $\text{GCD}(n, r - 1)$, the result must be p

- if p divides $r + 1$ then q cannot divide $r + 1$ and therefore q must divide $r - 1$. In this case $\text{GCD}(n, r - 1)$ yields q
- if for a given $a < n$ we find s such that $a^s = 1$, then with probability ≥ 0.5 we get $a^{s/2}$ as a nontrivial root of 1. Indeed:
 - by Chinese Remainder Theorem a number $a < n$ is represented by $a_p = a \bmod p$ and $a_q = a \bmod q$
 - given a and b we may compute representation of $a \cdot b \bmod n$ by computing $a_p \cdot b_p \bmod p$ and $a_q \cdot b_q \bmod q$
 - there are two roots of 1 modulo prime number p : 1 and $p - 1$
 - if $a^s = 1 \bmod n$, while $a^{s/2} \neq 1 \bmod n$, then $a^{s/2} \bmod p$ is 1 or -1
 - there are the following cases:
 1. $a^{s/2} = 1 \bmod p, a^{s/2} = -1 \bmod q$
 2. $a^{s/2} = -1 \bmod p, a^{s/2} = 1 \bmod q$
 3. $a^{s/2} = -1 \bmod p, a^{s/2} = -1 \bmod q$

the last case corresponds to $-1 \bmod n$, the first two ones to a nontrivial roots of -1
- so it suffices to find such an s - the order of a . By repeating the procedure for different a 's we finally find a nontrivial root of $-1 \bmod n$

Quantum operations and gates

- a quantum computer should perform some operations on qubits, technical realization is a challenge, but in theory possible
- we consider l - qubit numbers as representing numbers $\bmod 2^l$ (well, this is fuzzy as each bit is fuzzy as a qubit), in this way we get quantum state for each $a < q = 2^l$
- Hadamard transformation: an easy way to create a quantum state such that takes any value a (denoted $|a\rangle$) with the same probability. The way to achieve this is:
 - create the state $|0\dots 0\rangle$
 - apply Hadamard transformation gate to it
 - each coordinate is transformed by

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1, 1 \\ 1, -1 \end{pmatrix}$$

so $|0\rangle$ is transformed to

$$\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$

- Quantum Fourier transform:

- regular FT: (x_1, \dots, x_N) transformed to (y_1, \dots, y_N) where

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j \cdot e^{(2\pi i \cdot j \cdot k)/N}$$

- quantum:

$\sum x_i \cdot |i\rangle$ transformed to $\sum y_i \cdot |i\rangle$ where

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j \cdot e^{(2\pi i \cdot j \cdot k)/N}$$

- in other words:

$$|j\rangle \rightarrow \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{(2\pi i \cdot j \cdot k)/N} \cdot |k\rangle$$

„efficient implementation” based on similar algebra as for DFT

Shor’s algorithm (based on presentation of Eric Moorhouse)

1. fix q such that $2n^2 < q < 3n^2$, $q = 2^l$ (or a product of small primes)

we use states with $2l$ qubits, notation $|a, b\rangle$ or $|a\rangle|b\rangle$

2. prepare state $|0, 0\rangle$ and apply Hadamard transformation to the first register. Its result is a state

$$|\psi\rangle = \frac{1}{\sqrt{q}} \cdot \sum_{a=0}^{q-1} |a, 0\rangle$$

3. fix $x < n$ at random

4. to the state $|\psi\rangle$ apply the quantum transformation

$$|a, 0\rangle \rightarrow |a, x^a \bmod n\rangle$$

the result is

$$\frac{1}{\sqrt{q}} \cdot \sum_{a=0}^{q-1} |a, x^a \bmod n\rangle$$

(there is a theory how to make such a computation with quantum gates)

5. measure the second register. The result is some k . But then the measured state changes to

$$\frac{1}{\sqrt{M}} \cdot \sum_{a \in A}^{M-1} |a, k\rangle$$

where A is the set of all a such that $x^a = k \bmod n$

so $A = \{a_0, a_0 + r, a_0 + 2r, \dots\}$ where r is the order of x and $M = |A|$ (so $M \approx q/r$)

$$\frac{1}{\sqrt{M}} \cdot \sum_{d=0}^{M-1} |a_0 + d \cdot r, k\rangle$$

6. apply the DFT to the first register. This changes the state

$$\frac{1}{\sqrt{M}} \cdot \sum_{d=0}^{M-1} |a_0 + d \cdot r, k\rangle$$

to

$$\frac{1}{\sqrt{q \cdot M}} \cdot \sum_{c=0}^{q-1} \sum_{d=0}^{M-1} e^{2\pi i \cdot c(a_0 + d \cdot r)/q} \cdot |c, k\rangle$$

which is equal to

$$\sum_{c=0}^{q-1} \frac{e^{2\pi i \cdot c \cdot a_0/q}}{\sqrt{q \cdot M}} \sum_{d=0}^{M-1} e^{2\pi i \cdot c \cdot d \cdot r/q} \cdot |c, k\rangle$$

$$\sum_{c=0}^{q-1} \frac{e^{2\pi i \cdot c \cdot a_0/q}}{\sqrt{q \cdot M}} \sum_{d=0}^{M-1} \zeta^d \cdot |c, k\rangle$$

where

$$\zeta = e^{2\pi i \cdot c \cdot r/q}$$

7. measure the first register (this is the key moment!!)

- which c is read depends on the values of $\sum_{d=0}^{M-1} \zeta^d$ which in turn corresponds to the probability
- if $c \cdot r/q$ is not very close to an integer, then the sum is $\frac{1 - \zeta^M}{1 - \zeta}$
- if $c \cdot r/q$ is an integer, then we sum up M ones
- so the former case is unlikely and the readings are concentrated around values c such that

$$c/q \approx d/r$$

for an integer d

- the rest is a classical computation involving c, q and trying different d 's. The search space is relatively narrow

CATACRYPT catastrophe cryptography

- what happens if assumptions broken (e.g. DL solvable for some group)?
- "post-quantum crypto"

reality:

- post-quantum is at early stage, no industrial products, logistically impossible to replace
- no plans, scenarios, ...
- consequences of building a quantum computer or anything that breaks the current mechanisms:
 - i. option 1: the situation is well hidden, the party controlling the means uses it without leaving traces but still claims the system is secure

- ii. option 2: disclosed, call for massive conversion to new products (may occur even if there is no quantum computer but it works for increasing sale numbers)
 - iii. option 3: disclosed, neglected
 - catastrophe is already there
-

XI. COMMUNICATION SECURITY – SSL/TLS

mistakes:

- risk of common (standard) groups
- cryptanalysis: most efficient number field sieve (NFS):
 - complexity subexponential (for \mathbb{Z}_p it is

$$\exp(1.93 + o(1))(\log p)^{1/3}(\log \log p)^{2/3}$$
 - most time precomputation independent from the target number y (where $\log y$ to be computed in a given group)
 - the time dependant from y can be optimized to subexponential but much lower
 - 512-bit groups can be broken, MitM attack can be mounted
- standard safe primes – seem to be ok, but attacker can amortize the cost over many attacks
- TLS with DH: frequently “export-grade” DH with 512 bit primes, about 5% of servers support DHE_EXPORT, most servers (90% and more) use a few primes of a given length, after a precomputation breaking for a given prime: reported as 90 sec
- TLS: client wants DHE, server offers DHE_EXPORT, but one can manipulate the messages exchanged, so that the client treats the (p_{512}, g, g^b) as a response to DHE – it is not an implementation bug!
 - handshake time is a problem, but some protocols allow. sending TLS warning alert that reset the countdown
 - ephemeral key hashing
 - sometimes non safe prime used ($\frac{p-1}{2}$ composite), Pohling-Hellman method can be used
 - DH-768 breakable on academic level, DH-1024 on the state level
- recommendations:
 - avoid fixed prime groups
 - transition to EC
 - deliberately do not downgrade security even if seems to be ok

- follow the progress in computer algebra

Padding attack (Serge Vaudenay)

Scenario:

- for encryption the plaintext should have the length as a multiply of b
- always pad something
- if i positions have to be padded, then writes i 's there. de-padding is then easy.
- encrypt the resulting padded plaintext x_1, \dots, x_N in CBC mode with IV (fixed or random) and a block cipher Enc:

$$y_1 = \text{Enc}(\text{IV} \oplus a_1), \quad y_i = \text{Enc}(y_{i-1} \oplus x_i)$$

- properties of CBC:
 - efficiency
 - confidentiality limits: if IV fixed one can check that two plaintexts have the same prefix of a given size

attack:

- manipulate the ciphertext
- destination node decrypts, it can detect incorrect padding
- decision: what to do if the padding is incorrect? Each reaction is wrong:
 - reject: creates padding oracle (attacker tests the behavior)
 - proceed: enables manipulation of the plaintext data

last word oracle:

- goal: compute $\text{Dec}(y)$ for a block y
- create an input for padding oracle:
 - create a 2 block ciphertext: $r = r_1 \dots r_b$ chosen at random, $c := r|y$
 - oracle call: if $\text{Oracle}(c) = \text{valid}$, then $\text{Dec}(y) \otimes r$ should yield a correct padding. whp this happens if $y_b = r_b \oplus 1$
- it may happen that the oracle says valid because of other correct padding. The following procedure solves the problem (idea: change consecutive words in the padding until invalid:
 1. pick r_1, r_2, \dots, r_b at random, take $i = 0$
 2. put $r = r_1 r_2 \dots r_{b-1} (r_b \oplus i)$
 3. run padding oracle on $r|y$, if the result “invalid” then increment i and goto (2)

4. $r_b := r_b \oplus i$
5. for $j = b$ to 2:
 - $r := r_1 \dots r_{b-j} (r_{b-j+1} \oplus 1) r_{b-j} \dots r_b$
 - ask padding oracle for $r|y$, if “invalid” then output $(r_{b-j+1} \oplus j) \dots (r_b \oplus j)$ and halt
6. output $r_b \oplus 1$

block decryption oracle

let $a_1 \dots a_b$ be the plaintext of y

decryption:

- get a_b via the last word oracle
- proceed step by step learning a_{j-1} once a_j, \dots, a_b are already known
 1. set $r_k := a_k \oplus (b - j + 2)$ for $k = j, \dots, b$ /* preparing the values so that the padding values $(b - j + 2)$ appear at the end
 2. set r_1, \dots, r_{j-1} at random, $i := 0$ /* search for the value that makes a proper padding
 3. $r := r_1 \dots r_{j-2} (r_{j-1} \oplus i) r_j \dots r_b$
 4. if $O(r|y) = \text{invalid}$, then $i := i + 1$ and goto 3
 5. output $r_{j-1} \oplus i \oplus (b - j + 2)$

decryption oracle

- block by block
- the only problem with the first block if IV is secret

bomb oracles:

- padding oracle in SSL/TLS breaks the connection if a padding error occurs, so can be used only once
- bomb oracle: try a longer part at once

other paddings:

- $00 \dots 0n$ instead of $nn \dots n$ – also vulnerable
- $12 \dots n$ instead of $nn \dots n$ – also vulnerable

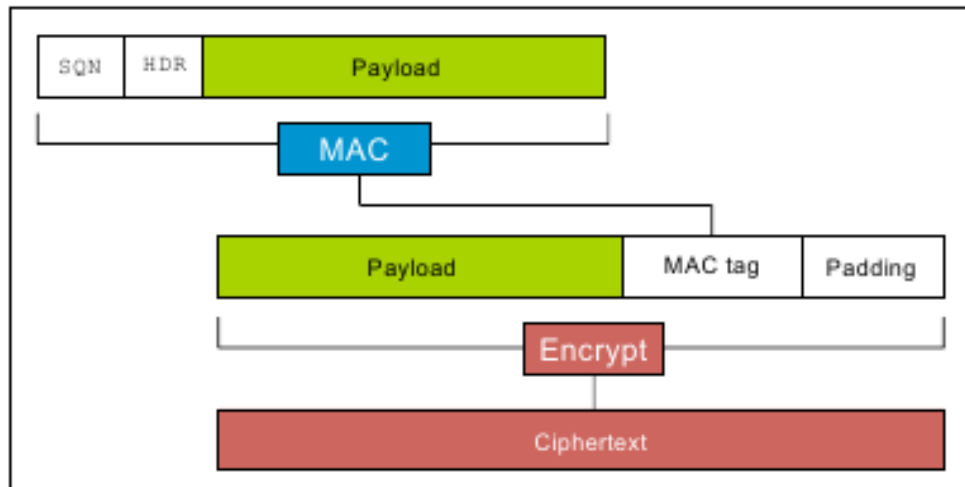
Applications for (old) versions of SSL/TLS, ...

- MAC applied before padding, so padding oracle techniques can be applied

- wrong MAC and wrong padding create the same error message - from SSL v3.0, debatable whether it is impossible to recognize situation via side channel (response time)
- TLS attempts to hide the plaintext length by variable padding:
 - checking the length of padding: take the last block y , send $r|y$ where the last word of r is $n \oplus 1$. acceptance means that the padding is of length n
 - checking paddings longer than a block: send ry_1y_2 where y_1y_2 are the last blocks
- IPSEC: discards message with a wrong padding, no error message, other activities to process errors (they may leak information)
- WTLS: decryption-failed message in clear (!) session not interrupted
- SSH: MAC after padding (+)

Lucky Thirteen

- concerns DTLS (similar to TLS for UDP connections)
- MAC-Encode-Encrypt paradigm (MEE), MAC is HMAC based



- 8-byte SQN, 5-byte HDR (2 byte version field, 1 byte type field, 2 byte length field)
- size of the MAC: 16 bytes (HMAC-MD5), 20 bytes (HMAC-SHA1), 32 bytes (HMAC-SHA-256)
- padding: $p + 1$ copies of p , at least one byte must be added
- after receiving: checking the details: padding, MAC, (underflow possible if padding manipulated and removing blindly)
- HMAC of M :

$$T = H((K_a \oplus \text{opad}) || H((K_a \oplus \text{ipad}) || M))$$

– **Distinguishing attack:**

- M_0 : 32 arbitrary bytes followed by 256 copies of 0xFF
- M_1 : 287 bytes followed by 0x00
- both 288 bytes, 18 plaintext blocks
- encoded $M_d||T||\text{pad}$, we aim to guess d
- C – the ciphertext
- create a ciphertext C' by truncating all parts corresponding to $T||\text{pad}$
- give $\text{HDR}||C'$ for decryption
- if M_0 : the 256 copies of 0xFF interpreted as padding and removed, remaining 32 bytes as short message and MAC, calculating MAC: 4 hash computed, then typically error returned to the attacker
- if M_1 : 8 hash evaluations

Plaintext recovery attacks

- C^* – the block of ciphertext to be broken, C' – the ciphertext block preceding it
- we look for P^* , where $P^* = \text{Dec}(C^*) \oplus C'$
- assume CBC with known IV, $b = 16$ (as for AES). $t = 20$ (as for HMAC-SHA-1)
- let Δ be a block of 16 bytes, consider

$$C^{\text{att}}(\Delta) = \text{HDR}||C_0||C_1||C_2||C' \oplus \Delta||C^*$$

it represents 4 non-IV blocks in the plaintext, the last block is:

$$P_4 = \text{Dec}(C^*) \oplus (C' \oplus \Delta) = P^* \oplus \Delta$$

- case 1: P_4 ends with 0x00 byte:
 - 1 byte of padding is removed, the next 20 bytes interpreted as MAC, 43 bytes left - say R . MAC computed on $\text{SQN}|\text{HDR}|R$ of 56 bytes
- case 2: P_4 ends with padding pattern of ≥ 2 bytes:
 - at least 2 bytes of padding removed, 20 bytes interpreted as MAC, at most 42 bytes left, MAC over at most $42+13=55$ bytes
- case 3: P_4 ends with no valid padding:
 - according to RFC of TLS 1.1, 1.2 treated as with no padding, 20 bytes treated as MAC, verification of MAC over $44+13=57$ bytes
 - MAC is computed to avoid other timing attack!

- time: case 1 and 3: 5 evaluations of SHA-1, case 2: 4 evaluations of SHA-1, detection of case 2 possible in LAN
- in case 2: most probable is the padding 0x01 0x01, all other paddings have probability about $\approx \frac{1}{256}$ of probability of 0x01 0x01, so we may assume that $P_4 = P^* \oplus \Delta$ ends with 0x01 0x01. Then we derive the last two bytes of P^* .
repeat the attack with Δ' that has the same last two bytes to check if the padding has the length bigger than 2.
- after recovery of the last two bytes the rest recovered byte by byte from right to left:
 - the original padding attack
 - e.g. to find 3rd rightmost byte set the last two bytes Δ so that P_4 ends with 0x02 0x02, then try different values for the Δ_{13} so that Case 2 occurs (meaning that P_4 ends with 3 bytes 0x02)
 - average time: $14 \cdot 2^7$ trials
- practical issues:
 - for TLS after each trial connection broken, so multi-session scenario
 - timing difference small, so necessary to gather statistical data
 - complexity in fact lower, since the plaintexts not from full domain: e.g. http username and password are encoded Base64
 - partial knowledge may speed up the recovery of the last 2 bytes
 - less efficient configuration of the lengths for HMAC-MD5 and HMAC-SHA-256

BEAST

attack, phase 0:

1. P to be recovered (e.g. a password, cookie, etc), requires ability to force Alice to put secret bits on certain positions
2. force Alice to send $0\dots 0P_0$ (requires malware on Alice computer) – of course encrypted
3. eavesdrop and get $C_p = \text{Enc}(C_{p-1} \oplus 0\dots 0P_0)$
4. guess a byte g
5. force Alice to send the plaintext $C_{i-1} \oplus C_{p-1} \oplus 0\dots 0g$
6. Alice sends $C_i = \text{Enc}(C_{i-1} \oplus C_{i-1} \oplus C_{p-1} \oplus 0\dots 0g) = \text{Enc}(C_{p-1} \oplus 0\dots 0g)$
7. if $C_i = C_p$ then $P_0 = g$

attack phase 1:

1. P_0 already known

2. force Alice to send $0...0P_0P_1$ and proceed as in phase 0

last phase: we get the test for the whole $P_0...P_{15}$

protection: browser must be carefully designed and do not admit injecting plaintexts (SOP- Same Origin Protection). Some products do not implement it.

CRIME (2012)

- based on compression algorithm used by some (more advanced) versions of TLS
- compression: LZ77 and then Huffman encoding, LZ77- sliding window approach: instead of a string put a reference to a previous occurrence of the same substring
- idea of recovering cookie:

```
POST / HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:14.0) Gecko/20100101 Firefox/3.6
Cookie: secretcookie=7xc89f94wa96fd7cb4cb0031ba249ca2
Accept-Language: en-US,en;q=0.8

(... body of the request ...)
```

Listing 1: *HTTP request of the client*

modified POST:

```
POST /secretcookie=0 HTTP/1.1
Host: example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:14.0) Gecko/20100101 Firefox/3.6
Cookie: secretcookie=7xc89f94wa96fd7cb4cb0031ba249ca2
Accept-Language: en-US,en;q=0.8

( ... body of the request ...)
```

Listing 2: *HTTP request modified by the attacker*

LZ77 compresses the 2nd occurrence of `secretcookie=` or `secretcookie=0`. We try all `secretcookie=i` to find out the case when compression is easier (`secretcookie=7`) when the first character recovered the attacker repeats the attack for the second character (trying all “`secretcookie=7i`” in the preamble)

TIME

- again based on compression but now on the server side (from the client to the server compression might be disabled and CRIME fails)
- works if the server includes the client’s request in the response (most do!)

- works even if SOP is enabled. SOP does not control data with the tag `img`, so the attacker can manipulate the length and therefore influence the number of blocks for block encryption
- the attacker requires malicious Javascript on the client's browser
- the attacker tries to get the secret value sent from the server to the client
- mechanism:
 - as in CRIME, the request sends "secretvalue=x" where x varies
 - the response is compressed, so it takes either "secretvalue=" or "secretvalue=x"
 - the length manipulated so that either two or one packets – connection specific data must be used: Maximum Transmission Unit
 - RTT (round trip time) measured
- independent on the browser, it is not an implementation attack!
- countermeasure: restrict displaying images

BREACH

Browser Reconnaissance and Exfiltration via Adaptive Compression of Hypertext

- attack against HTTP compression and not TLS compression as in case of CRIME
- a victim visits attacker-controlled website (phishing etc).
- force victim's computer to send multiple requests to the target website.
- check sizes of responses

```
GET /product/?id=12345&user=CSRFtoken=<guess> HTTP/1.1
Host: example.com
```

Listing 4: *Compromised HTTP request*

```
<form target="https://example.com:443/products/catalogue.aspx?id=12345&user=CSRFtoken=<guess>">
...
<td nowrap id="tdErrLgf">
<a href="logoff.aspx?CSRFtoken=4bd634cda846fd7cb4cb0031ba249ca2">Log Off</a></td>
```

Listing 5: *HTTP response*

- requirements: application supports http compression, user's input in the response, sensitive data in the response
- countermeasures:
 - disabling compression

- hiding length (randomizing the length of the output – it makes the attacks only harder if the attack can be repeated many times)
- no secrets in the same response as the user's data
- masking secret: instead of S send $R||S \oplus R$ for random R (fresh in each response)
- trace behaviour of requests and warn the user

POODLE (2014)

in SSL v.3.0 using technique from BEAST:

- padding is not covered by MAC so the attacker can manipulate it
- padding non-deterministic: padding 1 to L bytes (L = block length), the last byte denotes the number of preceding padding random bytes
- encrypted POST request:
POST /path Cookie: name=value... $\langle r \backslash n \backslash r \backslash n \rangle$ body ||20-byte MAC||padding
- manipulations such that:
 - the padding fills the entire block (encrypted to C_n)
 - the last unknown byte of the cookie appears as the last byte in an earlier block encrypted into C_i
- attack: replace C_n by C_i and forward to the server
usually reject
accept if $\text{Dec}_K(C_i)[15] \oplus C_{n-1}[15] = 15$, thereby $P_i[15] = 15 \oplus C_{n-1}[15] \oplus C_{i-1}[15]$
proceed in this way byte by byte
- downgrade dance: provoke lower level of protection by creating errors say in TLS 1.0, and create connection with SSL v3.0
- the attack does not work with weak (!) RC4 because of no padding

Weaknesses of RC4

- known weaknesses:
 - the first 257 bytes of encryption strongly biased, ≈ 200 bytes can be recovered if ≈ 232 encryptions of the same plaintext available
simply gather statistics as in case of Ceasar cipher
 - at some positions (multiplies of 256) if a zero occurs then the next position more likely to contain a zero
- broadcast attack: force the user to encrypt the same secret repeatedly and close to the beginning
- countermeasure: no secrets in the initial part!

TLS 1.2

differences with TLS 1.1 and TLS 1.0:

- explicit IV instead of implicit IV
 - IDEA and DES 64bit removed
 - MD5/SHA-1 PRF 65 is replaced with a suite specified hash function – SHA-256 for all TLS 1.2 suites, but in the future also SHA-3, ...
 - digitally-signed element includes the hash algorithm used
 - Verify_data length is no longer fixed length \Rightarrow TLS 1.2 can define SHA-256 based cipher suites
 - new encryption modes allowed: CCM, GCM
-

CCM encryption mode

Prerequisites: block cipher algorithm; key K ; counter generation function; formatting function; MAC length $Tlen$

Input: nonce N ; payload P of $Plen$ bits; valid associated data A

Computation: Steps:

1. formatting applied to (N, A, P) , result: blocks B_0, \dots, B_r
2. $Y_0 := \text{Enc}_K(B_0)$
3. for $i = 1$ to r : $Y_i := \text{Enc}_K(B_i \oplus Y_{i-1})$
4. $T := \text{MSB}_{Tlen}(Y_r)$
5. generate the counter blocks $\text{Ctr}_0, \text{Ctr}_1, \dots, \text{Ctr}_m$ for $m = Plen/128$
6. for $j = 0$ to m : $S_j := \text{Enc}_K(\text{Ctr}_j)$
7. $S := S_1 || \dots || S_m$
8. $C := (P \oplus \text{MSB}_{Plen}(S)) || (T \oplus S_0)$

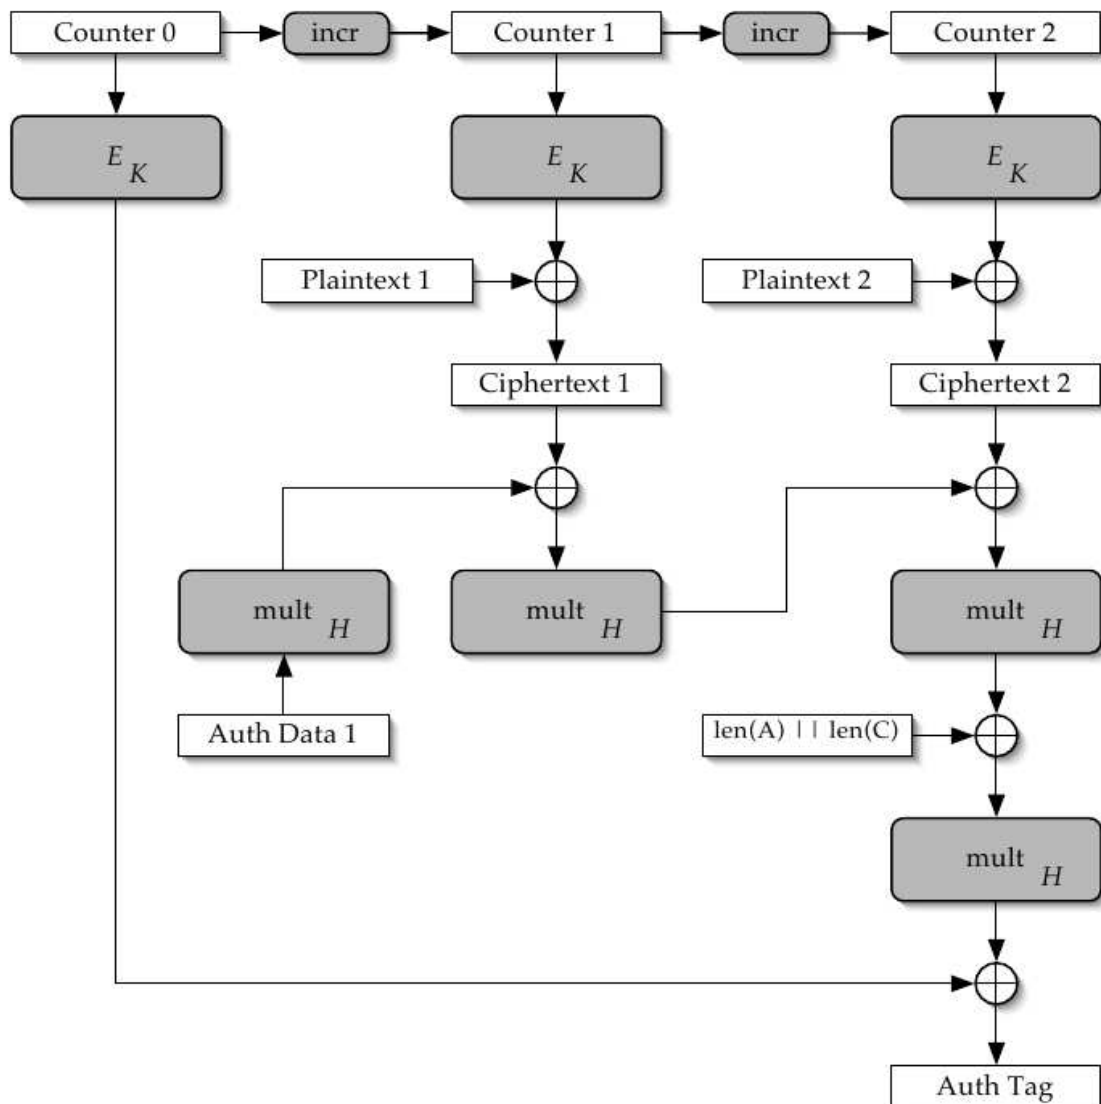
Decryption:

1. return INVALID, if $Clen < Tlen$
 2. generate the counter blocks $\text{Ctr}_0, \text{Ctr}_1, \dots, \text{Ctr}_m$ for $m = Plen/128$
 3. for $j = 0$ to m : $S_j := \text{Enc}_K(\text{Ctr}_j)$
 4. $S := S_1 || \dots || S_m$
 5. $P := \text{MSB}_{Clen}(C) \oplus \text{MSB}_{Plen}(S)$
 6. $T := \text{LSB}_{Tlen}(C) \oplus \text{MSB}_{Tlen}(S_0)$
 7. If N, A or P invalid, then return INVALID, else reconstruct B_0, \dots, B_r
 8. recompute Y_0, \dots, Y_r
 9. if $T \neq \text{MSB}_{Tlen}(Y_r)$, then return INVALID, else return P .
-

GCM (The Galois/Counter Mode)

Computation: Steps:

1. $H := \text{Enc}_K(0^{128})$
2. $Y_0 := \text{IV} || 0^{31}$ if length of IV should be 96
or $Y_0 := \text{GHASH}(H, \{\}, \text{IV})$
3. $Y_i := \text{incr}(Y_{i-1})$ for $i = 1, \dots, n$ (counter computation)
4. $C_i := P_i \oplus \text{Enc}_K(Y_i)$ for $i = 1, \dots, n-1$ (counter based encryption)
5. $C_n^* := P_n \oplus \text{MSB}_u(\text{Enc}_K(Y_n))$ (the last block need not to be full)
6. $T := \text{MSB}_t(\text{GHASH}(H, A, C)) \oplus \text{Enc}_K(Y_0)$



Details of computation of the tag

$\text{GHASH}(H, A, C) = X_{m+n+1}$ where m is the length of authenticating information A , and:

X_i equals:

$$0 \quad \text{for } i = 0$$

$$\begin{array}{ll}
(X_{i-1} \oplus A_i) \cdot H & \text{for } i = 1, \dots, m-1 \\
((X_{i-1} \oplus (A_m^* || 0^{128-v})) \cdot H & \text{for } i = m \\
(X_{i-1} \oplus C_i) \cdot H & \text{for } i = m+1, \dots, m+n-1 \\
((X_{m+n-1} \oplus (C_m^* || 0^{128-u})) \cdot H & \text{for } i = m+n \\
((X_{m+n} \oplus (\text{len}(A) || \text{len}(C))) \cdot H & \text{for } i = m+n+1
\end{array}$$

Decryption:

1. $H := \text{Enc}_K(0^{128})$
2. $Y_0 := \text{IV} || 0^{31}1$ if length of IV should be 96
or $Y_0 := \text{GHASH}(H, \{\}, \text{IV})$
3. $T' := \text{MSB}_t(\text{GHASH}(H, A, C)) \oplus \text{Enc}_K(Y_0)$, is $T = T'$?
4. $Y_i := \text{incr}(Y_{i-1})$ for $i = 1, \dots, n$
5. $P_i := C_i \oplus \text{Enc}_K(Y_i)$ for $i = 1, \dots, n$
6. $P_n^* := C_n^* \oplus \text{MSB}_u(\text{Enc}_K(Y_n))$

TLS 1.3 (draft)

- many old algorithms removed (RC4,static DH, ..)
- HKDF (hash key derivation function: first derive PRK via hashing of the shared secret+salt+user input, from PRK derive the secrets by hashing with sequence number)
- ChaCha20 stream cipher with Poly1305 authentication code
- Addition of the Ed25519 and Ed448 digital signature algorithms
- Addition of the x25519 and x448 key exchange protocols

ChCha

- stream cipher, Chacha extends 256 bit stream key into 2^{64} randomly accessible streams, each of 2^{64} block of 64 bytes
- Daniel Bernstein's construction,
- used by Google for communication between mobile devices and Google websites, also RFC
- variant of SALSA from European ECRYPT competition
- working on 4 32-bit words
- quarter-round of SALSA 20
 1. $b = b \text{ xor } (a + d) \lll 7$
 2. $c = c \text{ xor } (b + a) \lll 9$
 3. $d = a \text{ xor } (c + b) \lll 13$

4. $a = a \text{ xor } (d + c) \lll 18$
- quarter-round of ChaCha20 (better diffusion)
 1. $a = a + b$; $d = d \text{ xor } a$; $d = d \lll 16$
 2. $c = c + d$; $b = b \text{ xor } c$; $b = b \lll 12$
 3. $a = a + b$; $d = d \text{ xor } a$; $d = d \lll 8$
 4. $c = c + d$; $b = b \text{ xor } c$; $b = b \lll 7$
 - Chacha matrix 4x4: (input is block counter+nonce)

const	const	const	const
key	key	key	key
key	key	key	key
input	input	input	input
 - round: 8 quarter-rounds:
 - 1st column, 2nd column, 3rd column, 4th column
 - diagonals $\text{QUARTERROUND}(x_0, x_5, x_{10}, x_{15})$,
 $\text{QUARTERROUND}(x_1, x_6, x_{11}, x_{12})$
 $\text{QUARTERROUND}(x_2, x_7, x_8, x_{13})$
 $\text{QUARTERROUND}(x_3, x_4, x_9, x_{14})$
 - ChaCha20 - 20 rounds

Poly1305-AES

- 16 byte MAC, variable message length, 16 byte AES key, 16 byte additional key, 16 byte nonce
- the only way to break Poly is to break AES
- can be applied to any other cipher
- fast, no patents, ...
- Keys:
 - 16 bit AES key k
 - 16 bit key r with certain properties: $r = r_0 + r_1 + r_2 + r_3$ where
 $r_0 \in \{0, 1, \dots, 2^{28} - 1\}$, $r_1/2^{32} \in \{0, 4, 8, 12, \dots, 2^{28} - 4\}$,

Poly1035

- MAC algorithm
- works with AES, not weaker than AES, but if AES fails then reuse safely with a different encryption scheme
- low per message overhead (even for short messages)
- no long lookup tables, therefore fit into cache memory even if multiple keys used
- keys k - for AES, r - 16 byte, treated as little endian 128-bit number
- some limitations on r because of efficiency of implementation
 - $r = r_0 + r_1 + r_2 + r_3$ where
 - $r_0 \in \{0, 1, \dots, 2^{28} - 1\}$, $r_1/2^{32} \in \{0, 4, 8, 12, \dots, 2^{28} - 4\}, \dots$
- nonces: 16 bit, encrypted by AES
- message: divided into 16 byte chunks. Each chunk treated as 17-byte number with little-endian, where the most significant byte is an added 1, result: c_0, c_1, \dots, c_q
- authenticator

$$(((c_1 r^q + c_2 r^{q-1} + \dots + c_q r^1) \bmod 2^{130} - 5) + \text{AES}_k(\text{nonce})) \bmod 2^{128}$$

denoted also $H_r(m) + \text{AES}_k(\text{nonce})$

- $2^{130} - 5$ is a prime,
- a nonce must be used at most once
- security depends on the fact that for two messages m, m' of length L pbb that $H_r(m) = H_r(m') + g$ is at most $8\lceil L/16 \rceil / 2^{108}$ – small probabilities of differentials

XII. CERTIFICATES and – SSL/TLS

“Certified Lies”

- rogue certificates + MitM attack: the user believes that is directed elsewhere
- no control over root CA’s worldwide, indicated either by operating system or the browser
- compelled assistance from CA’s ?

ROGUE Certificates and MD5

- target: create a certificate (webserver, client) that has not been issued by CA

- not forging a signature contained in the certificate but:
 - find two messages that $\text{Hash}(M_0) = \text{Hash}(M_1)$ and M_0 as well as M_1 have some common prefix that you expect in a certificate (e.g. the CA name)
 - submit a request corresponding to M_0 , get a certificate with the signature over $\text{Hash}(M_0)$
 - copy the signature from the certificate concerning M_0 to a certificate based on M_1
- problems: some data in M_0 are to be guessed : sequential number, validity period, some other are known in advance: distinguished name, ...

legitimate website certificate		rogue CA certificate
serial number		serial number
issuing CA		issuing CA
validity period		validity period
domain name	chosen prefixes	rogue CA name
		1024 bit RSA public key
		extensions
		“CA=true”
		tumor
2048 RSA public key	collision bits	
extension “CA=false”	identical suffix	

Table.

- finding M_0 and M_1 has to be fast (otherwise the guess about the serial number and validity will fail) - e.g. a day over the weekend
- attack on MD5, general picture:

message A		message B
prefix P		prefix P'
padding S_r		padding S'_r
birthday blocks S_b		birthday blocks S'_b
near-collision block $S_{c,1}$		near-collision block $S'_{c,1}$
near-collision block $S_{c,2}$		near-collision block $S'_{c,2}$
...		...
near-collision block $S_{c,r}$	←collision→	near-collision block $S'_{c,r}$
suffix		suffix

Table.

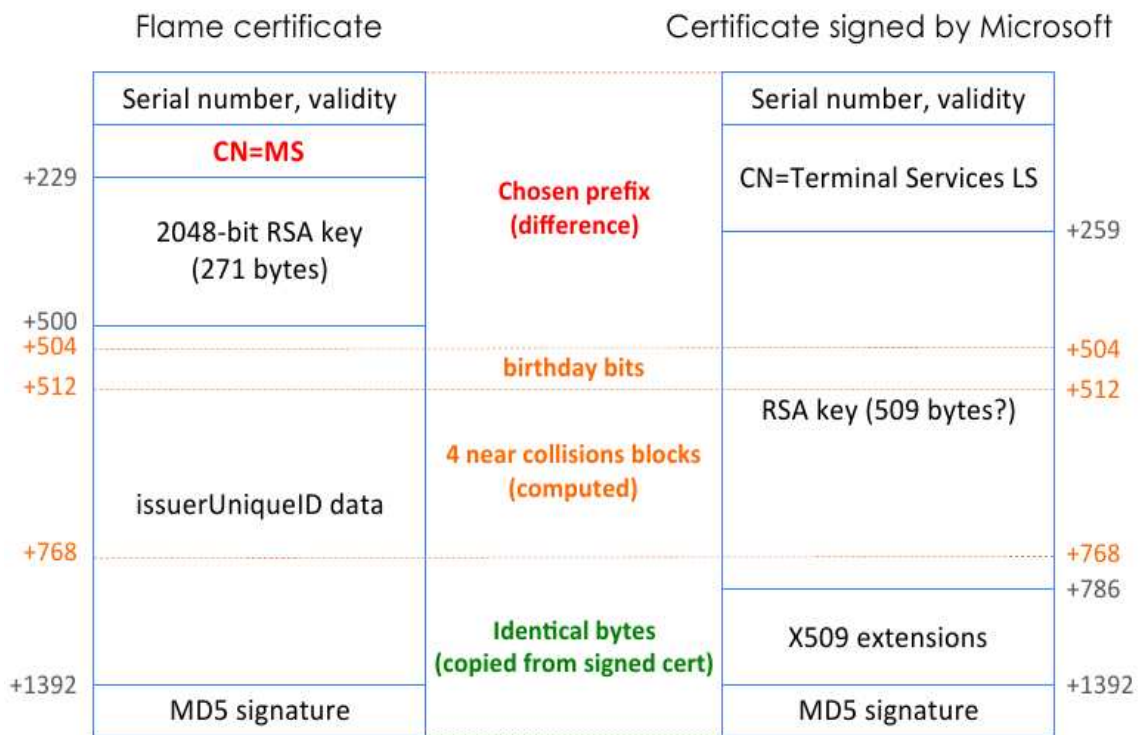
- identical prefix, birthday bits, near collision blocks:
 - birthday bits: 96, end at the block boundary, they are RSA bits – in the genuine certificate, “tumor” (ignored part by almost all software- marked as a comment extension) – in the rogue certificate
- birthday bits make the difference of intermediate hash values computed for both certificates fall into a *good class*

- then apply 3 near-collision blocks of 512-bits. website: we have “consumed” $208 + 96 + 3 \cdot 512 = 1840$ bits of the RSA modulus. Rogue certificate: all bits concerned are in the “tumor”
- after collision bits: $2048 - 1840 = 208$ bits needed to complete the RSA modulus of the webpage – we have to generate an RSA number with the prefix of 1840 bits already fixed
 - continued so that two prime factors:
 - B denotes the fixed 1840-bit part of the RSA modulus followed by 208 ones
 - select at random 224-bit integer q until $B \bmod q < 2^{208}$, continue until both q and $p = \lfloor B/q \rfloor$ are prime. Then
 - $p \cdot q$ is an RSA number
 - $p \cdot q < B$, $B - p \cdot q = B - q \cdot \lfloor B/q \rfloor < 2^{208}$. Hence $p \cdot q$ has the same 1840 most significant bits as B
 - this RSA number is not secure, but still factorizing it is not feasible and cannot be checked by CA before signing (as the smallest factor is more than 67-digit prime)
 - ... one can create RSA signature for the webpage for the certificate request
- attack complexity (number of hash block evaluations) for a chosen prefix MD5: 2^{49} at 2007, 2^{39} in 2009, not much motivation for more work - remove MD5 certificates! (For a collision: 2^{16})
for SHA-1 still 2^{77} in 2012 (for a collision: 2^{65})
- history:
 - attack found
 - real collision computed as a proof-of-concept
 - CA informed and given time
 - publication
 - code available

FLAME

- malware discovered 2012, 20MB, sophisticated code, mainly in Middle East, government servers attacked
- draft of the attack:
 - client attempts to resolve a computer name on the network, in particular make WPAD (Web Proxy Auto-Discovery Protocol) requests
 - Flame claims to be WPAD server, provides wpad.dat configuration file

- victim that gets wpad.dat sets its proxy server to a Flame computer (later no sniffing necessary!)
- Windows updates provided by the Flame computer. The main problem is that the updates must be properly signed!
- signatures obtained for terminal Services, certificates issued by Microsoft LSRA PA. No Extended Key Usage restrictions – allows code signing, (except for Microsoft Hydra X.509 extension – this cannot be used for code-signing on Vista and Windows 7)
- till 2012 still signatures with MD5 hash used
- MD5 collision necessary to remove extension



XIII. CACHE ATTACKS

idea:

- applies to multiprocess architectures, with strict separation between processes offered by the system: hypervisor and virtualization, sandboxing, ...
- trying to get secrets from one processes by another process with no privileges
- despite separation protection the processes share cache

- there is a strict control over the cache content but **cache hits and cache misses** might be detected by **timing for the attacker's process** (and not of the victim process)
- the timing for cache access should somehow depend on the sensitive information to be retrieved
- difficulty: other than in the classical cryptanalysis – access to plaintext or ciphertext might be impossible (they belong to the victim process) - the attacker can only predict something

cache:

- cache is necessary: gap between CPU speed and latency of memory access, innermost cache access $\approx 0.3\text{ns}$, main memory access $\approx 50\text{ns}$ to 150ns
- set-associative memory cache:
 - cache line of B bytes
 - a row consisting of W cache lines
 - S cache sets, each consisting of a line
 - when a cache miss occurs, then a memory block is copied into one of cache lines evicting its previous contents
 - a memory block with address a can be cached only into the cache set with the index i such that $i = \lfloor a/B \rfloor$ — **this is crucial for the attack**
- cache levels: slight complication to the attacks but differences of timing enable to recognize the situation

CASE STUDY: AES encryption

AES software implementation:

- particularly vulnerable because of its design
- AES defined in algebraic terms, but lookup table is typically faster
- there are arguments against algebraic implementations as the execution time may provide a side channel
- key expansion: round zero: simply the key bytes directly, other rounds: key expansion reversible (details irrelevant for the attack)
- fast implementation based on lookup tables T_0, T_1, T_2, T_3 and $T_0^{(10)}, T_1^{(10)}, T_2^{(10)}, T_3^{(10)}$ for the last round (with no MixColumns)
- round operation

$$\left(x_0^{(r+1)}, x_1^{(r+1)}, x_2^{(r+1)}, x_3^{(r+1)} \right) := T_0(x_0^r) \oplus T_1(x_5^r) \oplus T_2(x_{10}^r) \oplus T_3(x_{15}^r) \oplus K_0^{(r+1)}$$

$$\left(x_4^{(r+1)}, x_5^{(r+1)}, x_6^{(r+1)}, x_7^{(r+1)} \right) := T_0(x_4^r) \oplus T_1(x_9^r) \oplus T_2(x_{14}^r) \oplus T_3(x_3^r) \oplus K_1^{(r+1)}$$

$$\left(x_8^{(r+1)}, x_9^{(r+1)}, x_{10}^{(r+1)}, x_{11}^{(r+1)}\right) := T_0(x_8^r) \oplus T_1(x_{13}^r) \oplus T_2(x_2^r) \oplus T_3(x_7^r) \oplus K_2^{(r+1)}$$

$$\left(x_{12}^{(r+1)}, x_{13}^{(r+1)}, x_{14}^{(r+1)}, x_{15}^{(r+1)}\right) := T_0(x_{12}^r) \oplus T_1(x_1^r) \oplus T_2(x_6^r) \oplus T_3(x_{11}^r) \oplus K_3^{(r+1)}$$

attack notation:

- $\delta = B/\text{entrysize}$ of lookup table, typically: $\text{entrysize}=4\text{bytes}$, $\delta = 16$, (so δ entries of a lookup table are within the same cache line – this is a complication for the attack!)
- for a byte y let $\langle y \rangle = \lfloor y/\delta \rfloor$, it indicates a memory block of y in T_l
- if $\langle y \rangle = \langle z \rangle$ then x and y correspond to requests to the same memory block of the lookup table and therefore to the same cache line
- $Q_k(p, l, y) = 1$ iff AES encryption of plaintext p under key K accesses memory block of index y in T_l at least once in 10 rounds
- $M_k(p, l, y)$ a measurement that has expected value bigger in case when $Q_k(p, l, y) = 1$ then in case when $Q_k(p, l, y) = 0$

“synchronous attack”

- plaintext random but known, one can trigger encryption (e.g. for VPN with unknown key, dm-crypt of Linux)
- phase 1: measurements, phase 2: analysis
- from experiments: AES key recovered using 65 ms of measurements (800 writes) and 3 sec analysis
- **round-one attack:** the first round attacked
 - i. accessed indices are simply $x_i^{(0)} = p_i \oplus k_i$ for $i = 0, \dots, 15$
 - ii. finding information $\langle k_i \rangle$ of k_i – test candidates \bar{k}_i
 - iii. if $\langle k_i \rangle = \langle \bar{k}_i \rangle$ and $\langle y \rangle = \langle p_i \oplus \bar{k}_i \rangle$, then $Q_k(p, l, y) = 1$ for the lookup $T_l(x_i^{(0)})$
 - iv. if $\langle k_i \rangle \neq \langle \bar{k}_i \rangle$, then there is no lookup in block y for T_l during the first round, but
 - there are $4 \cdot 9 - 1 = 35$ other accesses affected by other plaintext bits during the entire encryption (4 per round, 9 rounds in total as the last round uses different look-up tables)
 - probability that none of them accesses block y for T_l is

$$\left(1 - \frac{\delta}{256}\right)^{35} \approx 0.104 \text{ for } \delta = 16$$
 - v. few dozens of samples required to find a right candidate for $\langle \bar{k}_i \rangle$
 - vi. together we determine $\log(256/\delta) = 4$ bits of each byte of the key
 - vii. no more possible for the first round, not enough to start brute force (still 64 bits to be found!)
 - viii. in reality more samples needed due to noise in measurements $M_k(p, l, y)$ and not $Q_k(p, l, y)$

- **two-round attack:** the second round attack because of the missing bits

i. exploiting equations derived from Rijndael specification:

$$x_2^{(1)} = s(p_0 \oplus k_0) \oplus s(p_5 \oplus k_5) \oplus 2 \bullet s(p_{10} \oplus k_{10}) \oplus 3 \bullet s(p_{15} \oplus k_{15}) \oplus s(k_{15}) \oplus k_2$$

$$x_5^{(1)} = s(p_4 \oplus k_4) \oplus 2 \bullet s(p_9 \oplus k_9) \oplus 3 \bullet s(p_{14} \oplus k_{14}) \oplus s(p_3 \oplus k_3) \oplus s(k_{14}) \oplus k_1 \oplus k_5$$

$$x_8^{(1)} = \dots$$

$$x_{15}^{(1)} = \dots$$

where s stands for the Rijndael Sbox, and \bullet means multiplication in the field with 256 elements

ii. lookup for $T_2(x_2^{(1)})$:

- $\langle k_0 \rangle, \langle k_5 \rangle, \langle k_{10} \rangle, \langle k_{15} \rangle, \langle k_2 \rangle$ already known
- low level bits of $\langle k_2 \rangle$ influence only low bits of $x_2^{(1)}$ so not important for cache access pattern
- the upper bits of $x_2^{(1)}$ can be determined after guessing low bits of k_0, k_5, k_{10}, k_{15} : there are δ^4 possibilities ($=16^4$)
- a correct guess yields a lookup in the right place
- an incorrect guess: some $k_i \neq \bar{k}_i$ so

$$x_2^{(1)} \oplus \bar{x}_2^{(1)} = c \bullet s(p_i \oplus k_i) \oplus c \bullet s(p_i \oplus \bar{k}_i) \oplus \dots$$

(for c depending on i) where ... depends on different random plaintext bits and therefore random

differential properties of AES studied for AES competition:

$$\Pr [c \bullet s(p_i \oplus k_i) \oplus c \bullet s(p_i \oplus \bar{k}_i) \neq z] > 1 - \left(1 - \frac{\delta}{256}\right)^3$$

so the false positive for lookup:

- $\left(1 - \frac{\delta}{256}\right)^3$ for computing $T_2(x_2^{(1)})$
- $\left(1 - \frac{\delta}{256}\right)$ for computing each of the remaining T_2
- together no access with pbb about $\left(1 - \frac{\delta}{256}\right)^{38}$
- this yields about 2056 samples necessary to eliminate all wrong candidates
- it has to be repeated 3 more times to get other nibbles of key bytes

iii. optimization: guess $\Delta = k_i \oplus k_j$ and take $p_i \oplus p_j = \Delta$, then i.e. $s(p_0 \oplus k_0) \oplus s(p_5 \oplus k_5)$ cancels out and we have to guess less bits (4 instead of 8)

- **similar attack: last round** - created ciphertext must be known to the attacker, otherwise similar. Subkey from the last round learnt, but key schedule is reversible

- **measurement: Evict+Time**
 - i. procedure:
 1. trigger encryption of p
 2. evict: access memory addresses so that one cache set overwritten completely
 3. trigger encryption of p
 - ii. in the evicted cache set one cache line from T_l
 - iii. measure time: if long, then cache miss and the encryption refers to evicted δ positions from the lookup table
 - iv. practical problem: triggering may invoke other activities and timing is not precise

- **measurement: Prime+Probe**
 - i. procedure
 1. (prime) read A : a contiguous memory of the size of the cache – results in overwriting the entire cache
 2. trigger an encryption of p (partial eviction at places where lookup used)
 3. (probe:) read memory addresses of A that correspond to $M_k(p, l, y)$
 - ii. easier: timing for probe suffice to check if encryption used a given cache set

- **complications in practice:**
 - i. adress of lookup tables in the memory - how they are loaded to the cache remains unknown – offset can be found by considering all offsets and then statistics for each offset (experiments show good results even on noisy environment)
 - ii. hardware prefetcher may disturb the effects. Solution: read and write the addresses of A according to a pseudorandom permutation

- **practical experiments:** e.g. Athlon 64, no knowledge of adresses mapping, 8000 encryptions with Prime & Probe
 Linux dm-crypt (disk, filesystem, file encryption): with knowledge of addressing, 800 encryptions (65 ms), 3 seconds analysis, full AES key

- **extensions of the attack:**
 - on some platforms timing shows also position of the cache line (better resolution for one-round attack)
 - remote attacks (VPN, IPSec): with requests that trigger immediate response (situation yet unclear about practicality)

“asynchronous attack”

- no knowledge of plaintext, no knowledge of ciphertext

- one-round attack
- based on frequency F of bytes in e.g. English texts, frequency score for each of $\frac{256}{\delta}$ blocks of length δ
- F is nonuniform: most bytes have high nibble equal to 6 (lowercase characters “a” through “o”)
- find j such that j is particularly frequent indicates $j = 6 \oplus \langle k_i \rangle$ and shows $\langle k_i \rangle$
- complication: this frequency concerns at the same time k_0, k_5, k_{10}, k_{15} affecting T_0 so we learn 4 nibbles but not their actual allocation to k_0, k_5, k_{10}, k_{15}
- the number of bits learnt is roughly: $4 \cdot (4 \cdot 4 - \log 4!) \approx 4 \cdot (16 - 3.17) \approx 51$ bits
- experiment: OpenSSL, measurements 1 minute, 45.27 bits of information on the 128-bit key gathered

Bernstein’s attack

- an alternative way of computing AES, algorithm applied in OpenSSL:
 - two constant 256-byte tables: S and S'
 - expanded to 1024-byte tables T_0, T_1, T_2, T_3

$$T_0[b] = (S'[b], S[b], S[b], S[b] \oplus S'[b])$$

$$T_1[b] = (S[b] \oplus S'[b], S'[b], S[b], S[b])$$

....
 - AES works with 16-byte arrays x and y , where x initialized with the key k , y initialized with $n \oplus k$, where n is the plaintext
 - AES computation is modifications of x and y :
 - x viewed as (x_0, x_1, x_2, x_3) (4 bytes parts)
 - $e := (S[x_3(1) \oplus 1], S[x_3(2)], S[x_3(3)], S[x_3(0)])$
 - replace (x_0, x_1, x_2, x_3) with $(e \oplus x_0, e \oplus x_0 \oplus x_1, e \oplus x_0 \oplus x_1 \oplus x_2, e \oplus x_1 \oplus x_2 \oplus x_3)$
 - modify y viewed as (y_0, y_1, y_2, y_3) , replace it with

$$(T_0[y_0[0]] \oplus T_1[y_1[1]] \oplus T_2[y_2[2]] \oplus T_3[y_3[3]] \oplus x_0,$$

$$(T_0[y_1[0]] \oplus T_1[y_2[1]] \oplus T_2[y_3[2]] \oplus T_3[y_0[3]] \oplus x_1,$$

$$(T_0[y_2[0]] \oplus T_1[y_3[1]] \oplus T_2[y_0[2]] \oplus T_3[y_1[3]] \oplus x_2,$$

$$(T_0[y_3[0]] \oplus T_1[y_0[1]] \oplus T_2[y_1[2]] \oplus T_3[y_2[3]] \oplus x_3$$
 - the next round uses $\oplus 2$ instead of $\oplus 1$ for x , otherwise the same. similar changes corresponding to rounds up to 9
 - in round 10 use $S[], S[], S[], S[]$ instead of T' 's
 - y is the final output

- embarrassing simple attack:
 - timing of execution depends on $k[i] \oplus n[i]$:
 - try many plaintexts
 - collect statistics for each byte as $n[i]$
 - the maximum occurs for z
 - the maximum corresponds to a fixed value for $k[i] \oplus n[13]$, say c
 - compute $k[13] = c \oplus z$
 - for different bytes different statistics observed: for some t a few values $k[t] \oplus \text{plaintext}[t]$, where substantially higher time observed
 - statistic gathered, different packet lengths
 - finally brute force checking all possibilities, nonce encrypted with the server key

Countermeasures

- "no reliable and practical countermeasure" so far
- implementation based on no-lookup but algebraic algorithm (slow!!!) or bitslice implementation (sometimes possible and nearly as efficient as lookup)
- alternative lookup tables: if smaller than less data leaks (but for cryptanalysis bigger Sboxes increase security)
- data-independent access to memory blocks - every lookup causes a redundant read in all memory blocks, generally: oblivious computation possible theoretically but overhead makes it impractical
- masking operations: \approx "we are not aware of any method that helps to resist our attack"
- cache state normalization: load all lookup tables - equires deep changes in OS and reduces efficiency, even then LRU cache policy may leak information which part has been used!
- process blocking: again, deep changes in OS
- disable cache sharing: deep degradation of performance
- "no-fill" mode during encryption:
 - preload lookup tables
 - activate "no-fill"
 - encrypt
 - deactivate "no-fill"

the first two steps critical and no other process is allowed to run

possible only in privileged mode, cost of operation prohibitive

- dynamic table storage: e.g. many copies of each table, or permute tables
details architecture dependent and might be costly
- hiding timing information: adding random values to timing makes the statistical analysis harder but still feasible
- try to protect some rounds (the first 2 and the last one) with any mean – but may be there are other attack techniques...
- cryptographic services at system level: good but unflexible
- sensitive status for user processes: erasing all data when interrupt
- specialized hardware support: seems to be the best choice
but the problem is not limited to AES or crypto – many sensitive data operations are not cryptographic and a coprocessor does not help

Meltdown – attack on modern processors

[please refer to the recent original paper](#)

XIV. ANONYMITY IN COMMUNICATION

- who is communicating with whom is a sensitive data

protection methods:

- broadcast channel
- token ring
- dining cryptographers -DC nets
- onion protocols and TOR

DC-nets

- protection of the message source. One of cryptographers is sending a 0 or 1.
- protocol for 3 cryptographers sitting at a round table:
 1. each pair of neighbors establish a shared bit at random
 2. each cryptographer that is not transmitting computes XOR of the shared bits that they knows,
 3. the sender computes the same XOR but swaps it if the bit transmitted is 1

4. each cryptographer reveals the computed result
5. every cryptographer computes XOR of all bits published

Steganography

- hiding transmission in innocent data (images, sound)
- procedure:
 1. original picture - stego image
 2. marking algorithm (maybe with a secret key) applied to message and stego image
 3. outcome transmitted/published
 4. retrieving the covered message
- invisibility property: without a key impossible to decide whether there is a message hidden
- typical applications: copyright protection, DRM,
- concrete techniques:
 - i. changing lsb bits of gray scale
 - ii. JPEG encoding: cosine transform, high frequency components are manipulated anyway for compression
 - iii. other digital transforms
 - iv. stretching (StirMark)
 - v. interference watermark
 - vi. echoing
 - vii. audio encoding: transformation and assigning coefficients to waves – manipulations of certain coefficients undetectable for listeners
 - viii. response times in transmission, packet order etc
- defense: (problematic or illegal due to intellectual property rights)
 - i. compression
 - ii. transforms and random distortions
 - iii. printing and scanning, input to an analog device and digitalize again

Onions

- limiting transmission overhead
- onions:
 - an onion created for a route going through servers A, B, C, \dots, Z to the final destination Γ

$$O_1 = \text{Enc}_A(B, \text{Enc}_B(C, \text{Enc}_C(\dots(\text{Enc}_Z(\Gamma, M))\dots)))$$

(by encryption Enc_X we mean encryption with the public key of server X)

- O_1 sent from the origin machine to A ,
- server A decrypts with its private key and gets B and $O_2 = \text{Enc}_B(C, \text{Enc}_C(\dots(\text{Enc}_Z(\Gamma, M))\dots))$
- A sends O_2 to B
- server B decrypts O_2 with its private key and gets C and $O_2 = \text{Enc}_C(\dots(\text{Enc}_Z(\Gamma, M))\dots)$
- the process is continued in the same way until server Z finds Γ, M and forwards the message M to machine Γ

each processing steps is like peeling off one layer of an onion

- idea: if two or more onion enter a server at the same time, get partially decrypted, then forwarded, then it is impossible to say which incoming onion corresponds to which outgoing onion - **node mixing**
- traffic analysis: assigning probabilities to permutations ($\pi(i) = j$ means that the i th sender has a message to the j th receiver)
- it is not enough to say that $\pi(i) = j$ with ppb $\approx \frac{1}{n}$ - example:
 - let us assume that the adversary knows that π is a circular shift
 - the adversary gets extra knowledge: if source i is talking to destination j , then $i + 1$ is talking to destination $j + 1$. However, still the property $\pi(i) = j$ is fulfilled
- question: necessary length of the onions? restricted case of n senders and n receivers, messages sent simultaneously
 - $O(\log^2 n)$ if adversary has a full knowledge of the system (not likely to have a better estimation unless ... big progress in math), assumption: uniform distribution for choosing destinations
 - $O(\log n)$ if the adversary can see only a constant fraction of nodes - tight bound, assumption: sender i may have non-uniform distribution of destination points

meaning of the results: *traffic analysis does not improve our prior knowledge in a significant way* (e.g. if we know in advance that source i always sends to destination j , then onions cannot hide this fact)

guarantees given in terms of total variation distance of two probability distributions:

$$\|\pi, \mu\| = \frac{1}{2} \sum_{\omega \in \Omega} |\pi(\omega) - \mu(\omega)|, \text{ where } \Omega \text{ is the set of all events}$$

Problems with onions

- replay attacks: just send the same onion (or partially decrypted onion) for the second time. the same subonions will appear along the forwarding path

defense: universal reencryption, example based on Elgamal:

- ciphertext of m :
 $(\beta^r, m \cdot g^r, \beta^s, g^s)$ for r, s chosen at random
- reencryption:
 1. choose α, β at random
 2. replace (y_1, y_2, y_3, y_4) by $(y_1 \cdot y_3^\alpha, y_2 \cdot y_4^\alpha, y_3^\beta, y_4^\beta)$
thereby we get $(\beta^{r+\alpha s}, m \cdot g^{r+\alpha s}, \beta^{s\beta}, g^{s\beta})$
if order of the group is a prime number then this is equivalent with choosing $(\beta^{r'}, m \cdot g^{r'}, \beta^{s'}, g^{s'})$ for random r', s'
- local view - not common lists of servers - long routes do not increase security:
 - give user A a list of servers with 50% of server used by nobody else
 - no matter how long is the routing path, it is likely that close to destination the path goes through a rogue server
 - a few destinations available from this rogue server (50% of cases the rogue one send directly to the destination)
 - an onion going through the rogue server originates from the attacked source
- timing at nodes - delays necessary
defense: collecting enough onions and flashing them at once. (slowdown!!!)
- sparse traffic = no protection

TOR

- free BSD licence
- connection based protocol, new connection established periodically ("10 minutes or so")
- routes limited to 3 TOR nodes
- onion based forwarding the symmetric keys
 - i. each node on the path learns only the predecessor and the successor
 - ii. the path established step by step:
 - after establishing a subpath X_0, X_1, \dots, X_k the subpath is used to send an encrypted message over the channel to X_k stating that the next node is X_{k+1} .
 - the sender and X_{k+1} negotiate a new connection key via DH key exchange
 - iii. after making a connection the message is encrypted symmetrically with the keys:
$$\text{AES}_{\text{relay1}}(\text{AES}_{\text{relay2}}(\text{AES}_{\text{relay3}}(m)))$$

each relay node removes one layer

- iv. the messages back to the sender: instead of decryption: encryption with keys shared with the sender. The sender has to decrypt the onion

problems:

- exit node knows the plaintext
 - traffic correlation
 - application level attacks
 - Heartbleed - change of public keys, some clients use old keys,
-

XV. PKI INFRASTRUCTURES

global versus local names

- OpenPGP
- X.509
- SPKI

anonymous authorization

- PKI
- BSI standards

X.509

- revocation: CRL, OCSP
- time limited validity (models: Austrian versus German to refresh validity of a document)

OpenPGP

- web of trust

flat infrastructure for biometric passports

- short time certificates for Document Verifier
- anchor points in passports for time

SPKI

- keys are principals for authentication
- basic name: key +identifier, key defines name space, identifier can refer to somebody in the name space
- compound names: key identifier1 identifier2 ... identifier n

- name certificate: $\text{Sing}_K(K, A, S, T)$, where K is the public key of name space, A is a basic name (no compound names), S is a subject (typically key or name), T is validity time
- no single certificate with given A , creating groups possible
- authorization certificates: $\text{Sign}_K(K, S, d, T, V)$ where K is the public key of name space, S is a local name or a key that receives authorization, d is delegation flag, T is a tag (interpretation subject to application), V is validity specification (time, condition, ...)
- chains of certificates, rewrite rules: each certificate is a directed edge: issuer gives the right to the subject, in fact transitive closure computed
- S-expressions - LISP like semantics for names, tags, ... There is logic to interpret multiple tags (intersection of rights)
- trust model unspecified – left for application level (X.509 explicit trust model)
- health certificate and suicide note

XVI. WIFI

standards:

- evolution
- little interaction with academic community
- underspecified,
- sometimes not literally implemented, lack of documentation

Krack

- attack based on crypto assumption: “no IV used twice”
- works despite “provable security” of WEP
- effects depend on particular implementation. Most cases:
 - decryption due to reuse of the same string in stream cipher
 - or just making mess by replay attack (e.g. against NTP- network time protocol)

[description in the original paper from 2017](#)

XVII. ACCESS CONTROL

(covered shortly in a different time)

ACCESS CONTROL

- decision whether a subject (user, etc) is allowed to carry out a specific action (operation) on an object (resource)
- policy – set of rules of (axioms + derivation rules) for making the decision
- main models:
 - Access Control Matrix or Access Matrix (AM)
 - Access Control List (ACL)
 - Role-Based Access Control (RBAC)
 - Attribute-Based Access Control (ABAC)
- principles:
 - least privilege
 - explicit permissions needed or admitted unless prohibited
 - rule administration is critical

AM

- Objects O, Subjects S, Access function
- Objects are protected
- a Subject may get access to an Object
- access function define concrete admissible operations (e.g. read,...)
- application areas: database systems, operating systems

problems of AM: lack of scalability, central administration, no flexibility, enormous effort

ACL

- object centric: for each object a list of admitted subjects
- corresponds to SPKI - a good cryptographic framework
- there are file systems based on ACL, Red Hat Linux: default ACL in a directory
- relational database systems, SQL systems, network administration

advantages of ACL: close to business models, simple, easy to implement in small and static systems, flexible,

problems of ACL:

- i. management of subjects is hard
- ii. poor scalability
- iii. changes might be tedious
- iv. impossible as global access control

alternative approach: *capability tickets* for subjects

RBAC

- organization:
 - subjects \leftrightarrow roles
 - roles \leftrightarrow permissions
- RBAC1: hierarchy of roles, a role gets all permissions from its lower roles in the hierarchy, roles form directed acyclic graph
- RBAC2: RBAC0+constraints (on mappings):
 - Mutually exclusive: a user can get only one role in a set, permission can be granted to only one role in a set, ...
 - cardinality: number of subjects with a given role
 - prerequisites: necessary to have another role before
 -
- RBAC3=RBAC1 consolidated with RBAC2
- RBAC sessions
- RBAC constraints
- advantages of RBAC:
 - simple management of users
 - manageable in large systems
- problems of RBAC:
 - no fine tuning of access rights
 - problems with dynamic changes
 - no useful in open systems

advantages of RBAC:

- simple management of users

- manageable in large systems

problems of RBAC:

- no fine tuning of access rights
- problems with dynamic changes
- no useful in open systems

ABAC

- components:
 - access control related entities
 - attributes of entities
 - policy
 - policy evaluation engine
- entities: requestor (Req), resource (Res), action (Act), environment (Env)
- attributes: anything (role, age, date and time (for Env)...))
- policy specification: logic of AC in terms of a policy
- policy evaluation: process answering access queries, computation based on entities, attributes and current policies – ADF (access decision function)
- standard: eXtensible Access Control Markup Language (XACML)
 - Policy Enforcement Point (PEP)
 - Policy Decision Point (PDP)
 - Policy Administration Point (PAP)
 - Policy Information Point (PIP)
- steps:
 - a request from the user gets to PEP
 - PEP sends the request to PDP
 - PDP gets attributes by querying PIP
 - PDP makes the decision and sends to PEP
 - PEP sends decision to the user
- major problems:
 - creating policies is a complicated task, errors are likely
 - trade-off between expressiveness and complexity of evaluation

CAAC (Context Aware Access Control)

- extension of RBAC
- context aware mapping: users-roles
- context aware mapping: roles-permissions

ICAM (Identity, Credentials Access Management)

RFC as an example of specification of a protocol

A. Exemplary RFC: draft of TLS 1.3 spec.

B. Required level of detail - ensuring unambiguous implementation.

C. Structure of the document: from general view to detailed description - to facilitate reading many datastructures and technicalities are shifted to the appendices.

- Abstract: What the document is about?
- Status: Internet draft, expiration date
- Copyright notice
- Table of contents:

1. Introduction

- a very nice note for RFC editors (present in the draft only)
- the goal of the protocol (authentication, confidentiality, integrity + definitions of the three terms, to let the non-cryptographers understand the document)
- the high level view - primary components
- conventions and terminology - for clear and precise understanding
- change log - for editors (present in the draft only)
- major differences from previous version of the protocol (TLS 1.2)

2. Protocol Overview. Handshake: what must be negotiated, what are the basic key exchange modes?

3. Presentation language: Big or little endian? basic block size? etc.

4.-9. Protocol components, further details.

10. Security considerations

D. Exemplary protocol detail: certificate request from server side (dnames of CAs) cross-certification.

VIII. SUBVERSION RESILIENCE

VIII. HOW TO CREATE A SYSTEM IN THE WORST POSSIBLE WAY?

Example: CHIP AUTHENTICATION PROGRAMME

“optimisation is the process of taking something that works and replacing it with something that almost works but is cheaper”

CAP - idea:

- cheap, Chip&PiN device
- keyboard, display, chip reader
- protecting PIN (it does not go to the PC)
- CAP device not personalized, one can use own card with CAP in a bank, or borrow from somebody
- recommended to use own CAP device
- optimized to be as much as possible based on EMV (standard for electronic purse)

used in UK, Master Card initiative

operation modes:

- identify: returns one-time code (like RSA-token) (based on symmetric key)
- respond: responds to a challenge using symmetric key
- sign: just as respond, however takes account number and value to generate the response

Protocol overview:

1. select application of the card (CAP has some fixed identifiers)
2. read records: account number, certificates , ... but important: CDOL1, CDOL2 (card object lists) and CAP (bit filter defining the protocol execution)
3. PIN verification
4. ciphertext generation: GENERATE AC command, response: Authorisation Request Cryptogram (ARQC), then the reader asks for Application Authentication Cryptogram (AAC) indicating cancelling the transaction (according to EMV)

challenge: AA (Authorized Amount), UN (Unpredictable number)

- for identify both are 0
- for respond: AA=0, UN=challenge

- for sign: AA=transaction value, UN= destination account

Response:

- based on the following data: ATC (application transaction counter), CID (Cryptogram Identification Data), IAD (Issuer Application Data - contains result of PIN verification), AC (Application Cryptogram - MAC (3DES CBC MAC) of the rest)
- CAP filter used to determine which bits to take
- NatWest: 5 least significant bits of ATC, 20 least significant bits of MAC, 1 bit from IAD
- Barclay: top bit of CID, 8 least significant bits of ATC, 17 least significant bits of MAC
- HBOS: top bit of CID, 7 least significant bits of ATC, 17 bits of MAC (not in one block), 1 bit from IAD

Verification: recomputed with the secret key shared with the card

Application:

- bank decides how to use (mode + semantic field)
- NatWest: respond mode, 8 bits of challenge, 4 random, 4 =last 4 digits of destination account, not used for login, transaction value not authenticated
- Barclay: identify necessary for login, for transaction: sign with destination account and transaction value (no freshness from bank, only ACT against replay – but might be played later)

Serious mistakes:

- checking PIN, result available on the device (mugging threat) – this concerns also cards of other banks
- the same PIN for ATM and online authentication – some keys on the CAP clean and some used - after stealing it one has 3 trials, 24 permutations on 4 keys, pbb to guess PIN to ATM becomes $\frac{1}{8}$
- CAP has no secret, infected PC may emulate CAP
- GSM in CAP to transmit secrets
- complicated instruction manual, the user may insert something else than intended account number
- overloading: sign with transaction with 0 value is valid for response (for a random account-nounce)
- NatWest: nonce as 4 digits in respond challenge, Chip&PIN terminal requests a number of responses from the card, later number of challenges from the online bank, there would be a match due to birthday paradox
(there are info indicating the attack: the number of requests, the change of transaction counter)

critical mistake: MITM regarding PIN verification

- PIN verification result never explicitly stated. Info to the bank contained in TVR (terminal verification results) and IAD (Issuer Application Data)
- TVR states possible failure conditions for authentication, in success not indicated which method used
 - bit8=1: carholder verification was not successful
 - bit7=1: unrecognized CVM
 - bit6=1: PIN Try limit exceeded
 - bit5=1: PIN required and PIN pad not present
 - bit4=1: PIN required, PIN pad present and PIN not entered
 - bit3=1: online PIN entered
- IAD may contain info on whether the PIN has been verified, but cannot be read by the terminal (proprietary format), So terminal can have a different picture of the situation
 - bit4=1: Issuer Authentication performed and failed
 - bit3=1: offline PIN performed
 - bit2=1: offline PIN verification performed and failed
 - bit1=1: unable to go online
- attack:
 1. tricking the terminal by sending 0x9000 to **Verify** without sending PIN to the card
 2. card thinks that the terminal is not supporting PIN and skip PIN or uses signature
 3. card does not increase PIN retry counter
 4. issuer thinks that the terminal was not supporting PIN and accepts
- practical case (as described in 2015 paper after 2011 case in Belgium)
 - credit cards stolen, used in Belgium, police used intersection analysis (card usage, SIM cards in the proximity) to identify the criminals
 - "minimal effort design", just to work. Implementation of the attack with MiTM
 - hardware: FUN chip attached to the original chip, wires connected (contacts of the FUN with contacts of the original chip), the card has traces of manipulation. thickness: .82 mm (instead of .76mm)
 - functional: data embossed on the card does not match the data from the chip, accepts any PIN, some wrong responses

What went wrong:

- no evaluation, no public certification report
- no reaction to S&P paper from 2011
- specification EMV: thousands of pages
- certification costs
- designing a solution: chaos, no sufficiently detailed documentation and regime
- CC very likely to fail:
 - asset: PIN, password, protected against use on a PC
 - no methodology to answer the question: **what are side-effects of protecting one asset**
 - important: security is **not monotonic**: improving situation with respect to one threat may worsen situation to another one. **Not reflected by CC framework.**
 - **optimization is necessary, but may lead to situation that is worse than the original one**

(other solution: a shadow PIN for the case of mugging)

IX. MODELLING UNSECURITY

attacks:

- hit-and-run
- hit-and-stay
- insider

life-cycle of a solution based on key secrecy:

1. T_0 : key created
2. T_1 : forensics-based key non-compromised
3. T_2 : key compromise (known to the adversary)
4. T_3 : forensics-based key already compromised
5. T_4 : key ceased from operation

$[T_1, T_4]$ is a gray period. Should be as short as possible

Example countermeasure: DSAS framework

archive for signatures,

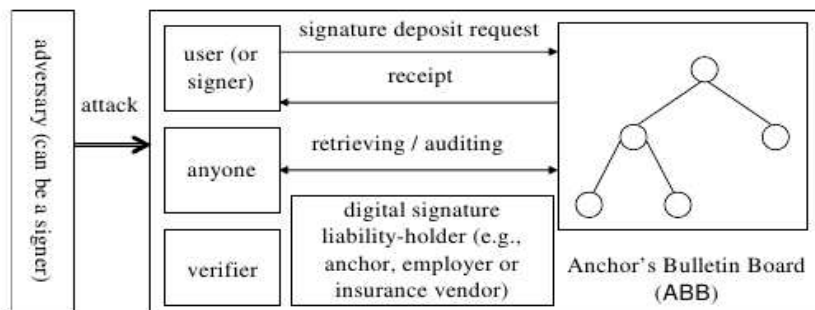




ABB:


- 3 levels of trees
- bottom level: signatures corresponding to one user in a separate tree, leftmost leaf holds public key, the root is a leaf in the level 2 tree
- middle level trees: binary tree for all user, the root is a leaf in top level tree
- top tree has nonleaf nodes corresponding to old roots and old leaves of trees of level 2




(a) ABB_0 at time t_0 (signatures-preserved case and signatures-compressed case): initialization with four user public keys



(b) ABB_1 at time t_1 (signatures-preserved case and signatures-compressed case): user pk_4 deposited two signatures



(c) ABB_2 at time t_2 (signatures-preserved case): user pk_4 deposited one new signature



(d) ABB_2 at time t_2 (signatures-compressed case): user pk_4 deposited one new signature

X. CRYPTOGRAPHIC FORENSICS

how to detect that a card has been cloned:

fail stop signatures: work if the clone created by cryptanalysis and deriving secret key and not by card inspection

- key generation by a trusted party: p, q chosen as for RSA, a - odd integer such that $\frac{p-1}{2a}$ is a prime and $q-1$ and a are coprime
- user chooses secret keys $sk_1, sk_2 \in Z_n^2$, public keys: $pk_1 = \alpha^a, pk_2 = sk_1^a \bmod n^2$
- signing m : $s := sk_1 \cdot sk_2^m$
- verification: $s^a := pk_1 \cdot pk_2^m$
- fail-stop idea: there are many solutions for sk_1 and sk_2 (namely a for each)
a different solution yields $s'^a = s^a \bmod n$, then $s'^a = s^a \bmod q, s' = s \bmod q$ since a is invertible modulo q

ctrl-signatures:

actors:

Inspection Authority: IA has a long period secret key k_{master} . For a user U , IA determines the control key $c_U := Hash_1(U, k_{master})$, and a pair of inspection keys: the private key $i_U = Hash_2(U, k_{master})$, and the public key $I_U = g^{i_U}$.

Card Issuer: for a user U , Card Issuer obtains the keys c_U and I_U from IA and installs them in the SSCD issued for U .

Signatories: the SSCD of a user U holds the preinstalled keys c_U and I_U , as well as the private signature key x_U created at random by the SSCD, and the public key $X_U = g^{x_U}$. (Note that the SSCD does not hold the key i_U .)

Certification Authority: CA has standard keys for issuing certificates for the public keys of the users, just as in PKI built according to the X.509 framework.

footprints:

Generating $f_U(k)$ - a hidden footprint for k and user U .

input: I_U, k

$f := Hash_3(I_U^k);$

output d least significant bits of f

For the inspection procedure carried out by Inspection Authority there is an alternative way for computing $f_U(k)$ (this is essential, since parameter k is present only on the SSCD):

Alternative generation of $f_U(k)$.

input: $i_U, r = g^k$

$f := Hash_3(r^{i_U});$

output d least significant bits of f

Creating the i th signature by SSCD of user U for message M .

input: a message M

“choose k at random so that $f_U(k) = \rho_U^i$ ”

proceed with the signing algorithm Sign with
the first signature component $r = g^k$

Inspection

INSPECTION OF A LIST OF SSCD SIGNATURES

Below we describe inspection of the signature list created by a user U .

1. User U presents a list S_1, S_2, \dots, S_t of allegedly all signatures created with SSCD of U , where the signatures appear on the list in the order in which they have been created. (If the signing time is included in the signatures, it is not necessary to specify the order of creating signatures.)
2. Apart from the regular verification of each signature S_i , the Inspection Authority checks all footprints. Namely, for each signature $S_j = (r_j, s_j), j \leq t$, IA computes the footprint $\omega_j := f_U(r_j)$.
3. If $(\omega_1, \omega_2, \dots, \omega_t) = (\rho_U^1, \rho_U^2, \dots, \rho_U^t)$, then inspection result is positive.